HP StorageWorks X9000 File Serving Software File System User Guide

Abstract

This guide describes how to configure and manage X9000 Software file systems and how to use NFS, CIFS, FTP, and HTTP to access file system data. The guide also describes the following file system features: quotas, remote replication, snapshots, data tiering, and file allocation. The guide is intended for system administrators managing X9300 Network Storage Gateway systems, X9320 Network Storage Systems, and X9720 Network Storage Systems.

HP Part Number: TA768-96035
Published: April 2011
Edition: Sixth

© Copyright 2009, 2011 Hewlett-Packard Development Company, L.P.

Confidential computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Acknowledgments

Microsoft, Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Revision History

Edition	Date	Software Version	Description
First	November 2009	5.3.1	Initial release of HP StorageWorks X9000 File Serving Software
Second	December 2009	5.3.2	Updated license and quotas information
Third	April 2010	5.4.0	Added information about file cloning, CIFS, directory tree quotas, the Statistics tool, and GUI procedures
Fourth	July 2010	5.4.1	Removed information about the Statistics tool
Fifth	December 2010	5.5.0	Added information about authentication, CIFS, FTP, HTTP, SSL certificates, and remote replication
Sixth	April 2011	5.6	Updated CIFS, FTP, HTTP, and snapshot information

Contents

1 Using X9000 Software file systems....8
File system organization and access....8
File system building blocks....9
Configuring file systems....9
Accessing file systems....10

2 Creating and mounting file systems....11
Creating a file system....11
Using 32-bit or 64-bit mode....11
Using the New Filesystem Wizard....11
Creating a file system using the CLI....14
Spillover files....15
Managing mountpoints and mount/unmount operations....15
GUI procedures....15
CLI procedures....16
Creating mountpoints....16
Deleting mountpoints....16
Viewing mountpoint information....16
Mounting a file system....16
Unmounting a file system....16
Mounting and unmounting file systems locally on Linux/Windows X9000 clients....17
Limiting file system access for X9000 clients....17
Using Export Control....18

3 Setting up quotas....19
How quotas work....19
Enabling quotas on a file system....19
Setting user and group quotas....19
Setting directory tree quotas....21
Using a quotas file....22
Importing quotas from a file....22
Exporting quotas to a file....22
Format of the quotas file....22
Using online quota check....23
Configuring email notifications for quota events....24
Deleting quotas....24
Deleting user and group quotas....24
Deleting directory tree quotas or usage limits....24

4 Maintaining file systems....25
Viewing information about file systems and components....25
Viewing physical volume information....25
Viewing volume group information....25
Viewing logical volume information....26
Viewing file system information....26
Lost+found directory....28
Viewing disk space information from a Linux X9000 client....29
Extending a file system....29
Rebalancing segments in a file system....30

How rebalancing works....30
Rebalancing segments on the management console GUI....30
Rebalancing segments from the CLI....31
Tracking the progress of a rebalance task....32
Viewing the status of rebalance tasks....32
Stopping rebalance tasks....32
Disabling 32-bit mode on a file system....32
Deleting file systems and file system components....33
Deleting a file system....33
Deleting segments, volume groups, and physical volumes....33
Deleting file serving nodes and X9000 clients....34
Checking and repairing file systems....34
Analyzing the integrity of a file system on all segments....34
Clearing the INFSCK flag on a file system....35
Troubleshooting file systems....35
ibrix_pv -a discovers too many or too few devices....35
Cannot mount on an X9000 client....35
NFS clients cannot access an exported file system....35
User quota usage data is not being updated....35
SegmentNotAvailable is reported....35
SegmentRejected is reported....36

5 Using NFS....38
Exporting a file system....38
Unexporting a file system....39
Autoconnecting NFS clients....40
Adding mount points to the autoconnect table....40
Deleting mount points from the autoconnect table....40
Setting up NFS clients....40

6 Configuring authentication for CIFS, FTP, and HTTP....42
Selecting an authentication method....42
X9000 management console requirement for Local Users authentication....44
Configuring local groups and local users....45
Configuring local group accounts....45
Configuring local users....47

7 Using CIFS....49
Configuring file serving nodes for CIFS....49
Starting or stopping the CIFS service and viewing CIFS statistics....49
CIFS shares....50
Managing CIFS shares with the X9000 management console GUI or CLI....50
Configuring SMB signing....51
Adding a CIFS share 52
Modifying a CIFS share....54
Deleting a CIFS share....54
Managing CIFS shares with Microsoft Management Console....54
Connecting to cluster nodes....54
Saving MMC settings....56
Granting share management privileges....56
Adding CIFS shares....57
Deleting CIFS shares....59
Linux static user mapping with Active Directory....59

Configuring Active Directory....59
Assigning attributes....61
Consolidating SMB servers with common share names....62
CIFS clients....63
Differences in locking behavior....64
Permissions in a cross-protocol CIFS environment....64
How the CIFS server handles UIDs and GIDs....64
Permissions, UIDs/GIDs, and ACLs....64
Pre-existing directories and files....65
New directories and files....65
Changing the way CIFS inherits permissions on files accessed from Linux applications....66
X9000 Windows Clients and Windows ACLs....66
Troubleshooting CIFS....66

8 Using FTP....68
Best practices for configuring FTP....68
Managing configuration profiles....69
Adding a configuration profile....69
Modifying a configuration profile....70
Viewing configuration profiles....70
Deleting a configuration profile....70
Managing FTP shares....70
Adding an FTP share....70
Modifying an FTP share....72
Viewing FTP shares....72
Deleting an FTP share....72
The vsftpd service....72
Starting or stopping the FTP service manually....73
Accessing shares....73

9 Using HTTP....75
Best practices for configuring HTTP....75
Configuring HTTP with the HTTP Wizard....76
Managing configuration profiles....78
Adding a configuration profile from the CLI....78
Modifying a configuration profile....78
Viewing configuration profiles....78
Deleting a configuration profile....78
Managing virtual hosts....78
Adding a virtual host....78
Modifying a virtual host....79
Viewing a virtual host....79
Deleting a virtual host....79
Managing HTTP shares....79
Adding an HTTP share....80
Modifying an HTTP share....80
Viewing HTTP shares....80
Deleting an HTTP share....80
Starting or stopping the HTTP service manually....80
Accessing shares....81

10 Managing SSL certificates....83
Creating an SSL certificate....83

Adding a certificate to the cluster....85
Exporting a certificate....86
Deleting a certificate....86

11 Using remote replication....87
Overview....87
Continuous or run-once replication modes....87
Remote cluster or intracluster....88
Many-to-many or many-to-one replications....88
Configuring the target for remote replication....89
GUI procedure....89
CLI procedure....90
Registering source and target clusters....90
Exporting the target file system....90
Identifying host and NIC preferences on the target cluster....91
Configuring and managing replication tasks on the GUI....91
Configuring and managing replication tasks from the CLI....93
Starting a remote replication task to a remote cluster....93
Starting an intra-cluster remote replication task....93
Starting a run-once directory replication task....93
Stopping a remote replication task....93
Pausing a remote replication task....93
Resuming a remote replication task....93
Querying remote replication tasks....94
Viewing replication status and activity....94
Configuring remote failover/failback....94
Troubleshooting remote replication....94

12 Creating snapshots....96
Setting up snapshots....96
Preparing the snapshot partition....96
Registering for snapshots....97
Discovering LUNs in the array....97
Reviewing snapshot storage allocation....97
Automated snapshots....97
Creating automated snapshots using the GUI....98
Creating automated snapshots from the CLI....101
Creating a snapshot scheme....101
Scheduling and starting automated snapshots....102
Other automated snapshot procedures....102
Modifying an automated snapshot scheme....103
Viewing automated snapshot schemes....103
Deleting an automated snapshot scheme....103
Managing snapshots....103
Creating a snapshot....103
Mounting a snapshot....103
Recovering system resources on snapshot failure....103
Deleting snapshots....103
Viewing snapshot information....104
Listing snapshot information for all hosts....104
Listing detailed information about snapshots....104
Accessing snapshot file systems....105
Troubleshooting snapshots....107

13 Using data tiering....108
Overview....108
Moving files between tiers....108
Writing a rule to implement a policy....108
Tiered file systems....109
Creating a file system that uses data tiering....109
Expanding a file system....110
Allocation policy....110
Managing a tiered file system and tiering policy....111
Assigning segments to tiers....111
Removing tier assignments....111
Deleting a tier....111
Listing tier information....111
Listing tiering policy information....111
Deleting a tiering policy rule....112
Starting and stopping a tiering operation....112
Reviewing tiering job status....112
Writing tiering rules....112
Operators and date/time qualifiers....112
Rule keywords....113
Migration rule examples....114
Ambiguous rules....114

14 Using file allocation....116
Overview....116
File allocation policies....116
How file allocation settings are evaluated....117
Using hostgroups for file allocation settings....117
When file allocation settings take effect on X9000 clients....117
Guidelines for using file allocation CLI commands....117
Setting file and directory allocation policies....117
Setting a file allocation policy....118
Setting a directory allocation policy....118
Setting segment preferences....118
Creating a pool of preferred segments....118
Restoring the default segment preference....119
Tuning allocation policy settings....119
Listing allocation policies....119

15 Support and other resources....120
Contacting HP....120
Related information....120
HP websites....120
Subscription service....120

Glossary....121

Index....123

1 Using X9000 Software file systems

File system organization and access

The following diagram shows how data is organized on a file system and how it is accessed.

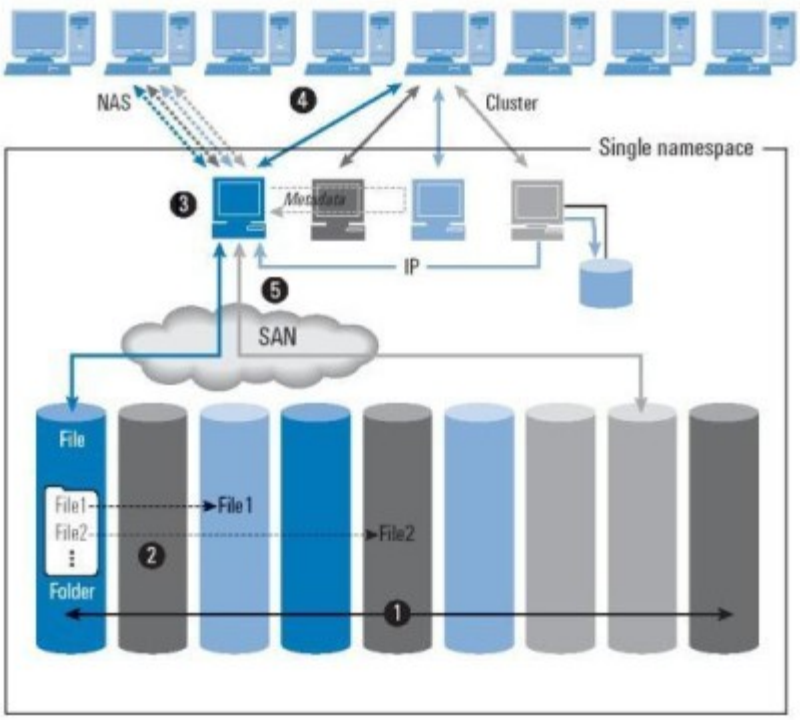

The diagram includes the following items:

1. The file system is a collection of segments (logical volumes) that organize data for faster access. Each segment is a repository for files and directories with no implicit namespace relationships among them. (A segment need not be a complete, rooted directory tree.) Segments can be of any size, and different segments can be of different sizes. A file can span several segments and multiple segments can be accessed in parallel within the same namespace.
2. The location of files and directories within segments is independent of their physical locations. A directory can be located on one segment, while the files in that directory are spread over other segments. Segments for new files and directories are selected dynamically according to an allocation policy. This policy is set by the system administrator based on anticipated access patterns and criteria such as performance and manageability.
3. File serving nodes, or servers, manage the individual segments of the file system. Each segment is assigned to one server, and each server can "own" multiple segments, as shown by the colors in the diagram. Segment ownership can be migrated from one server to another while the file system is in use. When servers are added to the cluster, the ownership of existing segments is distributed for proper load balancing and utilization by all servers. When storage is added, ownership of the new segments is distributed among existing servers.
4. Clients run the applications that use the file system. Clients can access the file system either as a locally mounted cluster file system using the X9000 client driver, or by using standard NAS protocols such as NFS, CIFS, HTTP, and FTP. Based on the file or directory being accessed, X9000 client requests are routed directly to the correct node. A client using NAS

protocols must mount the file system from a file serving node. All requests are sent to the mounting server, which performs the required routing.

5. A client request can be made for a file on a segment that is either owned by the server, owned by another server but accessible by this server over the SAN, or owned by another server and not accessible by this server over the SAN. In the second scenario, the server obtains the relevant metadata from the owning server and performs the I/O directly over the SAN. In the third scenario, the I/O is performed through the owning server over the IP network.

File system building blocks

A file system is created from building blocks. The first block comprises the underlying physical volumes, which are combined in volume groups. Segments (logical volumes) are created from the volume groups. The built-in volume manager handles all space allocation considerations involved in file system creation.

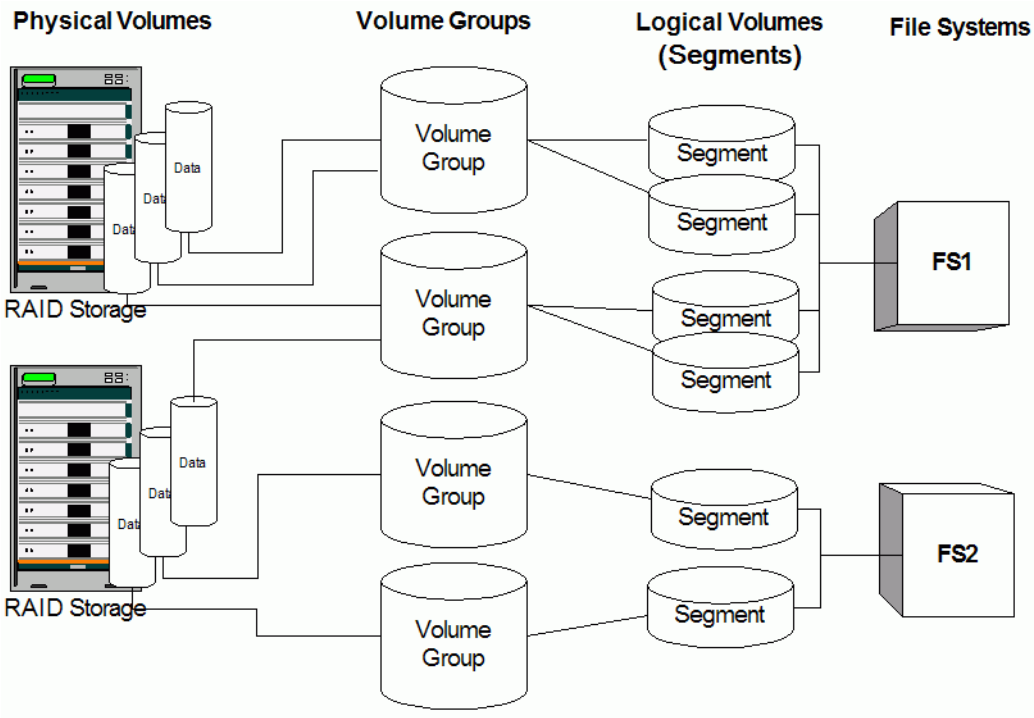

Configuring file systems

You can configure your file systems to use the following features:

- **Quotas.** This feature allows you to assign quotas to individual users or groups, or to a directory tree. Individual quotas limit the amount of storage or the number of files that a user or group can use in a file system. Directory tree quotas limit the amount of storage and the number of files that can be created on a file system located at a specific directory tree. See "Setting up quotas" (page 19).
- **Remote replication.** This feature provides a transparent method to replicate changes in a source file system on one cluster to a target file system on either the same cluster or a second cluster. See "Using remote replication" (page 87).
- **Snapshots.** This feature allows you to capture a point-in-time copy of a file system for online backup purposes and to simplify recovery of files from accidental deletion. The snapshot replicates all file system entities at the time of capture and is managed exactly like any other file system. See "Creating snapshots" (page 96).
- **Data tiering.** This feature allows you to set a preferred tier where newly created files will be stored. You can then create a tiering policy to move files from initial storage, based on file

attributes such as such as modification time, access time, file size, or file type. See "Using data tiering" (page 108).

- **File allocation.** This feature allocates new files and directories to segments according to the allocation policy and segment preferences that are in effect for a client. An allocation policy is an algorithm that determines the segments that are selected when clients write to a file system. See "Using file allocation" (page 116).

Accessing file systems

Clients can use the following standard NAS protocols to access file system data:

- **NFS.** See "Using NFS" (page 38) or more information.
- **CIFS.** See "Using CIFS" (page 49) for more information.
- **FTP.** See "Using FTP" (page 68) for more information.
- **HTTP.** See "Using HTTP" (page 75) for more information.

You can also use X9000 clients to access file systems. Typically, these clients are installed during the initial system setup. See the *HP StorageWorks X9000 File Serving Software Installation Guide* for more information.

2 Creating and mounting file systems

This chapter describes how to create file systems and mount or unmount them.

Creating a file system

You can create a file system using the New Filesystem Wizard provided with the management console GUI, or you can use CLI commands. The New Filesystem Wizard also allows you to create an NFS export or a CIFS share for the file system.

Using 32-bit or 64-bit mode

A file system can be created to use either 32-bit or 64-bit mode. In 32-bit mode, clients can run both 32-bit and 64-bit applications. In 64-bit mode, clients can run only 64-bit applications. If all file system clients (NFS, CIFS, and X9000 clients) will run only 64-bit applications, HP recommends that you enable 64-bit mode because more inodes will be available per segment for the applications.

File systems created with 32-bit mode compatibility can be converted later to allow clients to run 64-bit applications (see "Disabling 32-bit mode on a file system" (page 32)). This is a one-time-only operation and cannot be reversed. If clients may need to run a 32-bit application, do not disable 32-bit mode.

Using the New Filesystem Wizard

To start the wizard, click **New** on the Filesystems page. The wizard includes six steps and a summary, starting with selecting the storage for the file system. For details about completing each step, see the GUI online help.

On the Select Storage dialog box, select the storage that will be used for the file system. The total size of the storage is displayed at the bottom of the screen. If your cluster includes storage that has not yet been discovered by the X9000 software, click **Discover**.

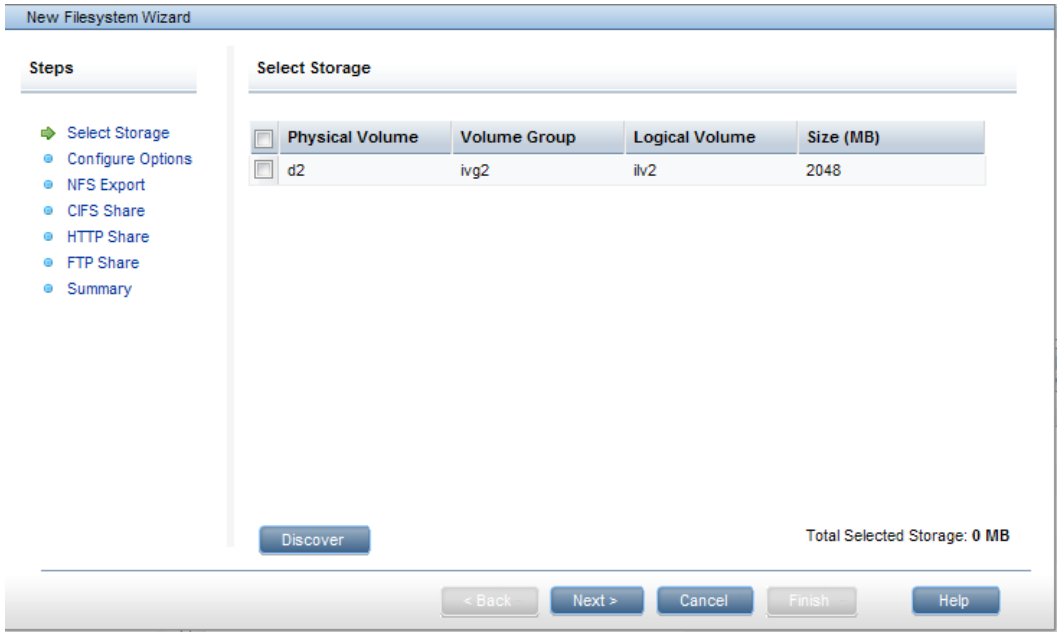

On the Configure Options dialog box, supply a name for the file system, and specify the appropriate configuration options.

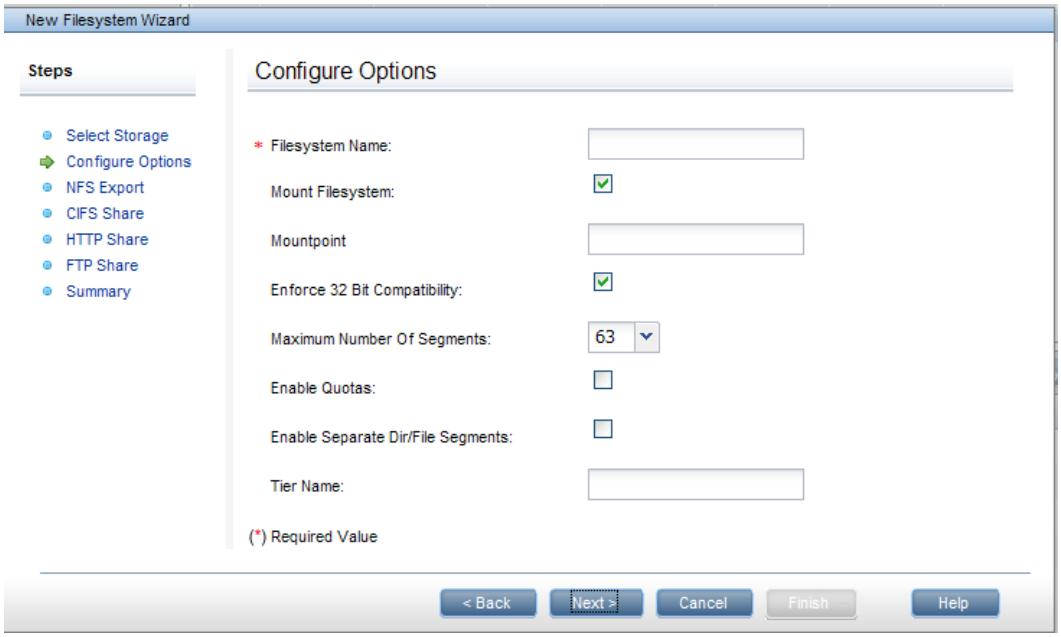

If the file system will be exported through NFS, configure an NFS export record on the NFS Export dialog box.

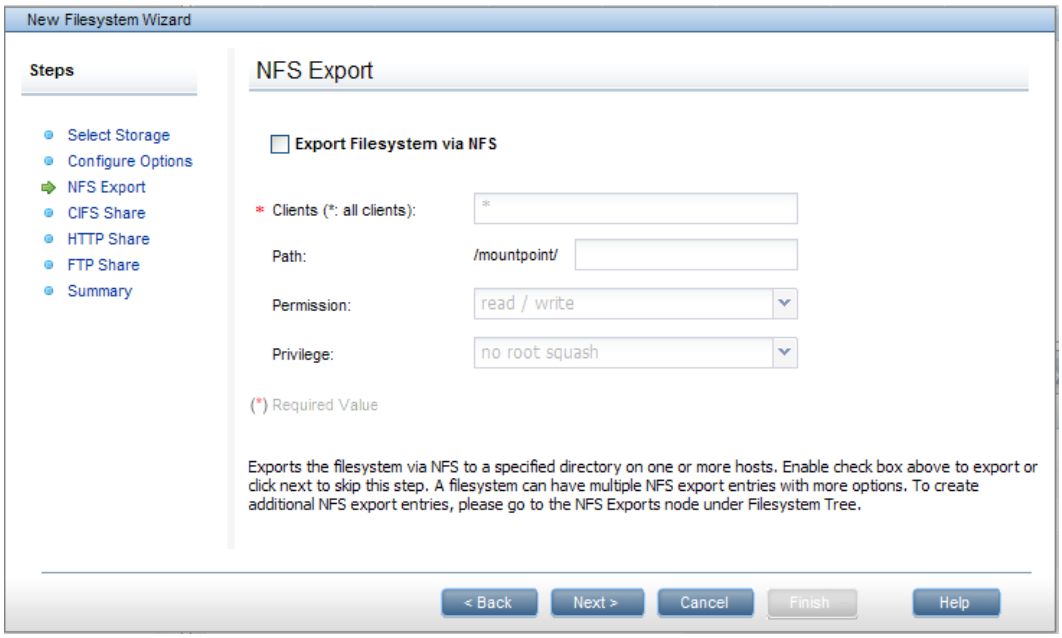
New Filesystem Wizard
Steps
Select Storage
Configure Options
NFS Export
CIFS Share
HTTP Share
FTP Share
Summary
NFS Export
Export Filesystem via NFS
* Clients (*: all clients):
*
Path:
/mountpoint/
Permission:
read / write
Privilege:
no root squash
(*) Required Value
Exports the filesystem via NFS to a specified directory on one or more hosts. Enable check box above to export or click next to skip this step. A filesystem can have multiple NFS export entries with more options. To create additional NFS export entries, please go to the NFS Exports node under Filesystem Tree.
< Back
Next >
Cancel
Finish
Help

If the file system will be made available to Windows clients through a CIFS share, create a share on the CIFS Share dialog box. For more information, see "Using CIFS" (page 49).

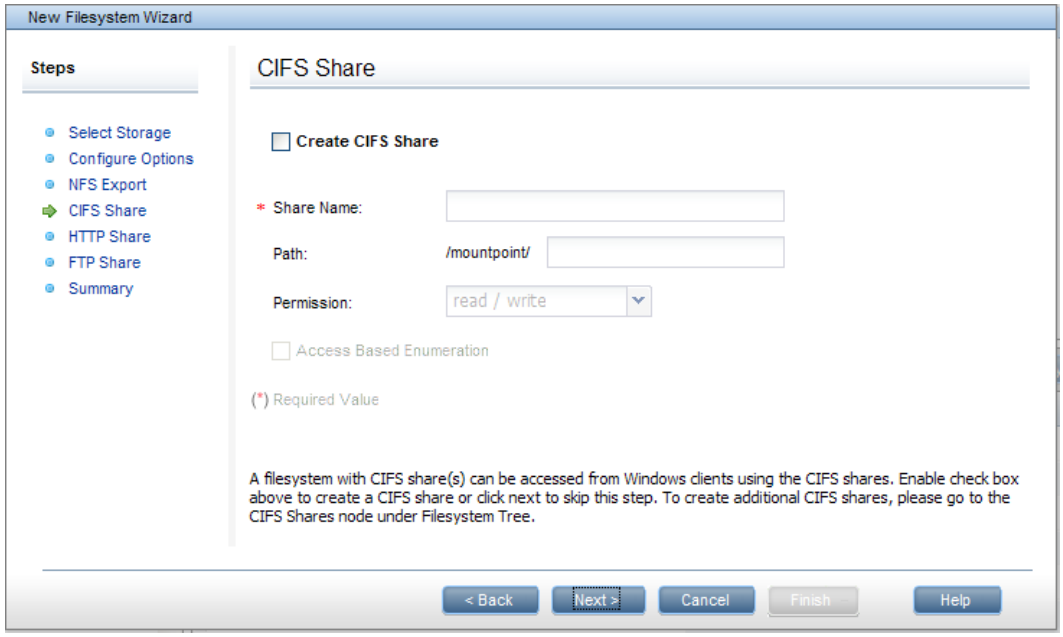

If clients will access the file system using HTTP or HTTPS, create an HTTP share on the HTTP Export dialog box. An HTTP configuration profile and an HTTP Vhost must already exist. If the Directory Path includes a subdirectory, be sure to create the subdirectory on the file system and assign read/write/execute permissions to it. (X9000 Software does not create the subdirectory if it does not exist, and instead adds a `/pub/` directory to the share path.) For more information, see "Using HTTP" (page 75).

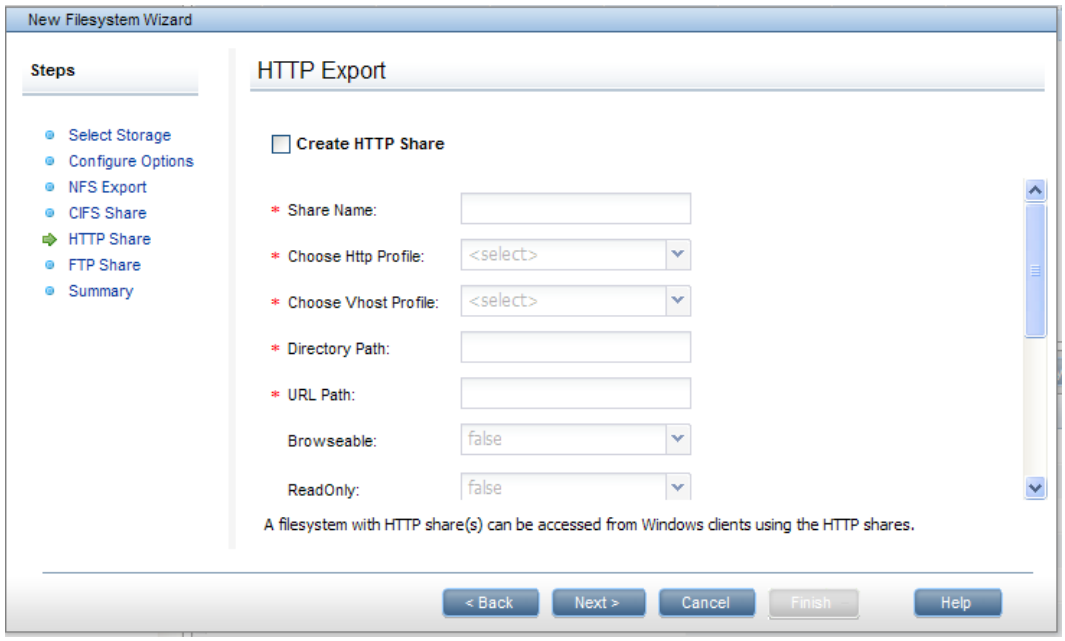

If clients will access the share using FTP or FTPS, create an FTP share on the FTP Share dialog box. An FTP configuration profile and an SSL certificate, if used, must already exist. If the Directory Path includes a subdirectory, be sure to create the subdirectory on the file system and assign read/write/execute permissions to it. (X9000 Software does not create the subdirectory if it does not exist, and instead adds a `/pub/` directory to the share path.) For more information, see "Using FTP" (page 68).

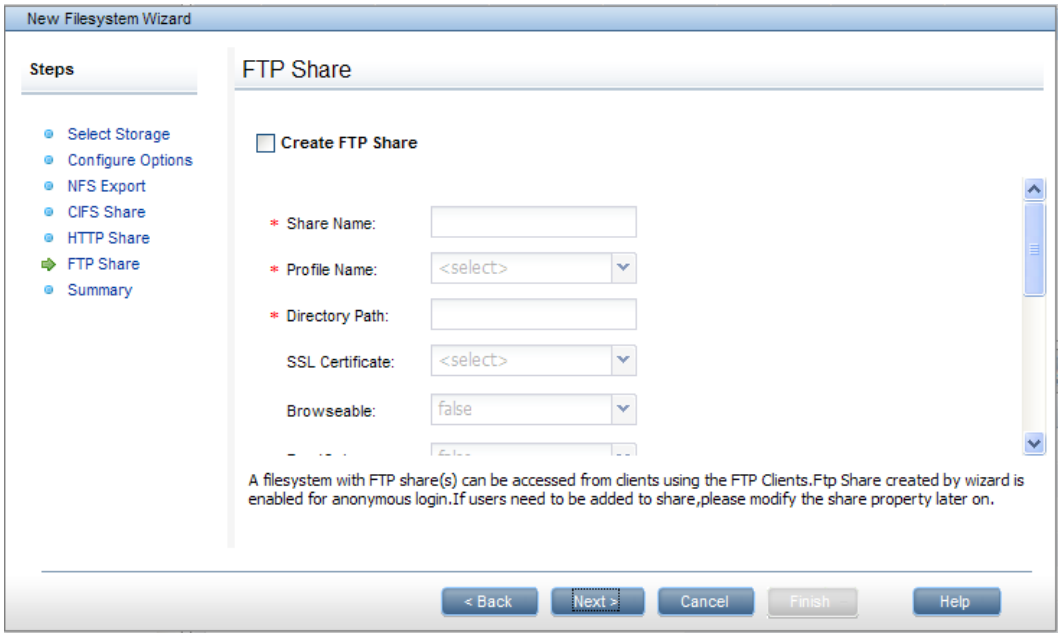

Review the Summary to ensure that the file system is configured properly. If necessary, you can return to a dialog box and make any corrections.

Creating a file system using the CLI

The `ibrix_fs` command is used to create a file system. It can be used in the following ways:

- Create a file system with the specified segments (segments are logical volumes):

  ```
  ibrix_fs -c -f FSNAME -s LVLIST [-t TIERNAME] [-a] [-q] [-o
  OPTION1=VALUE1,OPTION2=VALUE2,...] [-t TIERNAME] [-F
  FMIN:FMAX:FTOTAL] [-D DMIN:DMAX:DTOTAL]
  ```

- Create a file system and assign specify segments to specific file serving nodes:

  ```
  ibrix_fs -c -f FSNAME -S LV1:HOSTNAME1,LV2:HOSTNAME2,... [-a] [-q]
  [-o OPTION1=VALUE1,OPTION2=VALUE2,...] [-t TIERNAME] [-F
  FMIN:FMAX:FTOTAL] [-D DMIN:DMAX:DTOTAL]
  ```

- Create a file system from physical volumes in a single step:

  ```
  ibrix_fs -c -f FSNAME -p PVLIST [-a] [-q] [-o
  OPTION1=VALUE1,OPTION2=VALUE2,...] [-t TIERNAME] [-F
  FMIN:FMAX:FTOTAL] [-D DMIN:DMAX:DTOTAL]
  ```

Creating a file system manually from physical volumes

This procedure is equivalent to using `ibrix_fs` to create a file system from physical volumes in a single step. Instead of a single command, you build the file system components individually:

1. Discover the physical volumes in the system. Use the `ibrix_pv` command.
2. Create volume groups from the discovered physical volumes. Use the `ibrix_vg` command.
3. Create logical volumes (also called segments) from volume groups. Use the `ibrix_lv` command.
4. Create the file system from the new logical volumes. Use the `ibrix_fs` command.

See the *HP StorageWorks X9000 File Serving Software CLI Reference Guide* for details about these commands.

Spillover files

The X9000 Software file system supports spilling over sequentially written files from one segment to another.

Managing mountpoints and mount/unmount operations

GUI procedures

When you use the New Filesystem Wizard to create a file system, you can specify a name for the mountpoint and indicate whether the file system should be mounted after it is created. The wizard will create the mountpoint if necessary. Click **Mount** or **Unmount** as necessary to mount or unmount the file system.

To view the mountpoint information for a file system, select the file system on the GUI Filesystems page, and click **Mountpoints** in the lower Navigator. The Mountpoints section shows the hosts that have mounted the file system, the name of the mountpoint, the access (RW or RO) allowed to the host, and whether the file system is mounted.

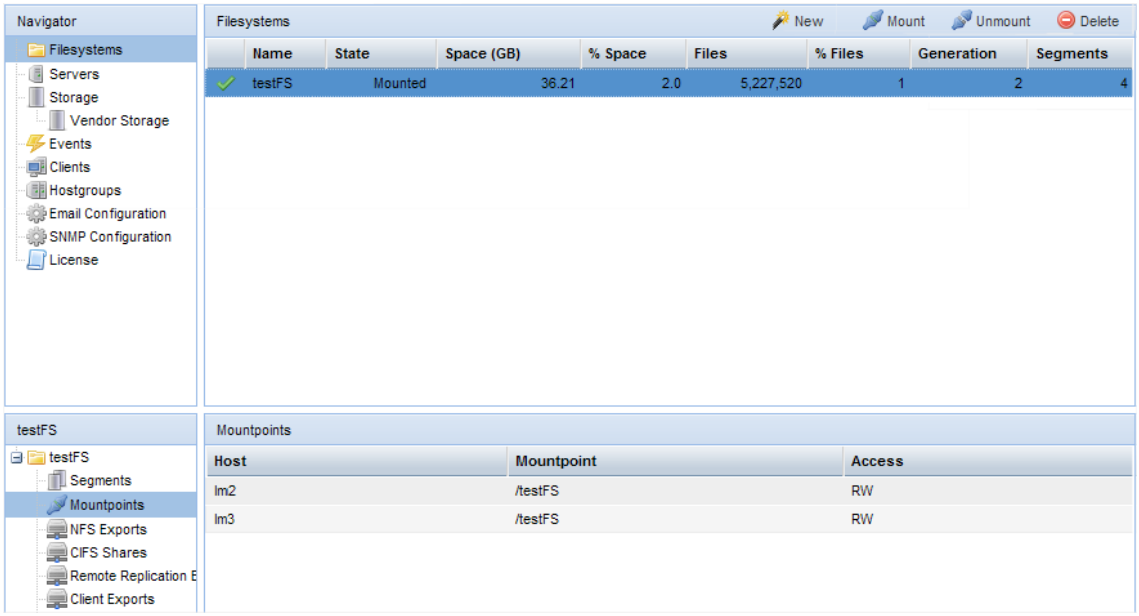

You can also view mountpoint information for a particular server. Select that server on the Servers page, and then select Mountpoints from the lower Navigator. To delete a mountpoint, select that mountpoint and click **Delete**.

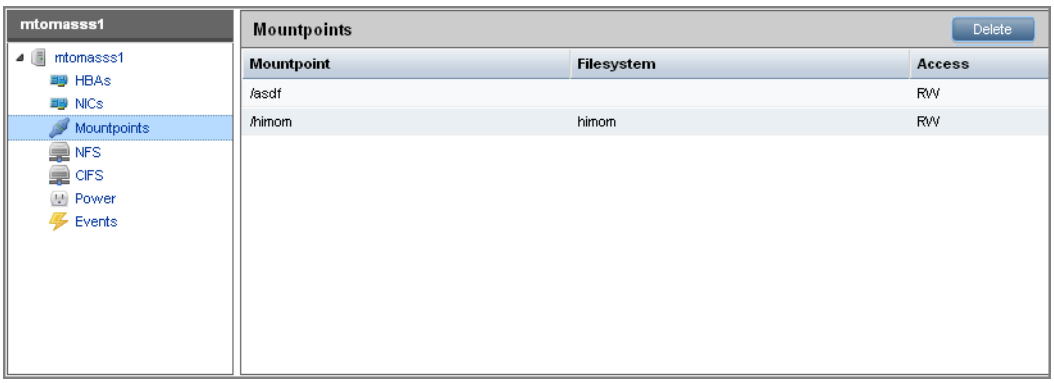

CLI procedures

The CLI commands are executed immediately on file serving nodes. For X9000 clients, the command intention is stored in the management console. When X9000 Software services start on a client, the client queries the management console for any commands. If the services are already running, you can force the client to query the management console by executing either `ibrix_client` or `ibrix_lwmount -a` on the client, or by rebooting the client.

If you have configured hostgroups for your X9000 clients, you can apply a command to a specific hostgroup. For information about creating hostgroups, see the administration guide for your system.

Creating mountpoints

Mountpoints must exist before a file system can be mounted. To create a mountpoint on file serving nodes and X9000 clients, use the following command.

```
<installdirectory>/bin/ibrix_mountpoint -c [-h HOSTLIST] -m MOUNTPOINT
```

To create a mountpoint on a hostgroup , use the following command:

```
<installdirectory>/bin/ibrix_mountpoint -c -g GROUPLIST -m MOUNTPOINT
```

Deleting mountpoints

Before deleting mountpoints, verify that no file systems are mounted on them. To delete a mountpoint from file serving nodes and X9000 clients, use the following command:

```
<installdirectory>/bin/ibrix_mountpoint -d [-h HOSTLIST] -m MOUNTPOINT
```

To delete a mountpoint from specific hostgroups, use the following command:

```
<installdirectory>/bin/ibrix_mountpoint -d -g GROUPLIST -m MOUNTPOINT
```

Viewing mountpoint information

To view mounted file systems and their mountpoints on all nodes, use the following command:

```
<installdirectory>/bin/ibrix_mountpoint -l
```

Mounting a file system

File system mounts are managed with the `ibrix_mount` command. The command options and the default file system access allowed for X9000 clients depend on whether the optional Export Control feature has been enabled on the file system (see "Using Export Control" (page 18) for more information). This section assumes that Export Control is not enabled, which is the default.

NOTE: A file system must be mounted on the file serving node that owns the root segment (that is, segment 1) before it can be mounted on any other host. X9000 Software automatically mounts a file system on the root segment when you mount it on all file serving nodes in the cluster. The mountpoints must already exist.

Mount a file system on file serving nodes and X9000 clients:

```
<installdirectory>/bin/ibrix_mount -f FSNAME [-o {RW|RO}] [-O MOUNTOPTIONS] -h HOSTLIST
-m MOUNTPOINT
```

Mount a file system on a hostgroup:

```
<installdirectory>/bin/ibrix_mount -f FSNAME [-o {RW|RO}] -g GROUP -m MOUNTPOINT
```

Unmounting a file system

Use the following commands to unmount a file system.

NOTE: Be sure to unmount the root segment last. Attempting to unmount it while other segments are still mounted will result in failure. If the file system was exported using NFS, you must unexport it before you can unmount it (see "Exporting a file system" (page 38)).

To unmount a file system from one or more file serving nodes, X9000 clients, or hostgroups:

```
<installdirectory>/bin/ibrix_umount -f FSNAME [-h HOSTLIST | -g GROUPLIST]
```

To unmount a file system from a specific mountpoint on a file serving node, X9000 client, or hostgroup:

```
<installdirectory>/bin/ibrix_umount -m MOUNTPOINT  [-h HOSTLIST | -g GROUPLIST]
```

Mounting and unmounting file systems locally on Linux/Windows X9000 clients

On both Linux and Windows X9000 clients, you can locally override a mount done on the management console. For example, if the configuration database on the management console has a file system marked as mounted for a particular client, that client can locally unmount the file system, thus overriding the management console.

Linux X9000 clients

To mount a file system locally, use the following command on the Linux X9000 client. A management console name (*fmname*) is required only if this X9000 client is registered with multiple management consoles.

```
ibrix_lwmount -f [fmname:]fsname -m mountpoint [-o options]
```

To unmount a file system locally, use one of the following commands on the Linux X9000 client. The first command detaches the specified file system from the client. The second command detaches the file system that is mounted on the specified mountpoint.

```
<installdirectory>/bin/ibrix_lwumount -f [fmname:]FSNAME
```

```
<installdirectory>/bin/ibrix_lwumount -m MOUNTPOINT
```

Windows X9000 clients

Use the Windows X9000 client GUI to mount file systems locally. Click the **Mount** tab on the GUI and select the cluster name from the list (the cluster name is the management console name). Then, enter the name of the file system, select a drive, and click **Mount**.

If you are using Remote Desktop to access the client and the drive letter is not displayed, log out and log back in. This is a known limitation of Windows Terminal Services when exposing new drives.

To unmount a file system on the Windows X9000 client GUI, click the **Umount** tab, select the file system, and then click **Umount**.

Limiting file system access for X9000 clients

By default, all X9000 clients can mount a file system after a mountpoint has been created. To limit access to specific X9000 clients, create an access entry. When an access entry is in place for a file system (or a subdirectory of the file system), it enters secure mode, and mount access is restricted to clients specified in the access entry. All other clients are denied mount access.

Select the file system on the Filesystems page, and then select **Client Exports** in the lower navigator. On the Create Client Export(s) dialog box, select the clients or hostgroups that will be allowed access to the file system or a subdirectory of the file system.

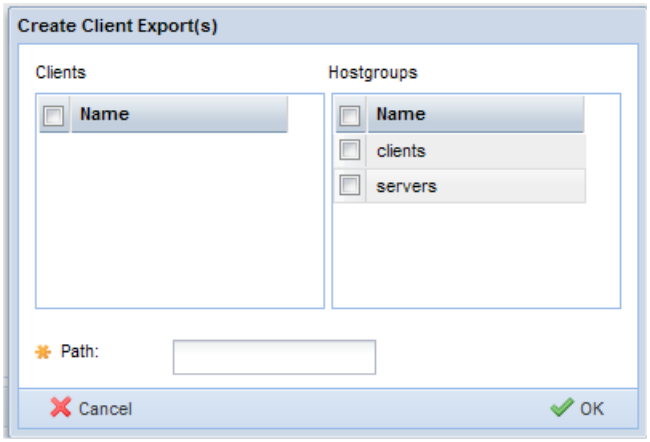

To remove a client access entry, select the affected file system on the GUI, and then select **Client Exports** from the lower Navigator. Select the access entry from the Client Exports display, and click **Delete**.

On the CLI, use the `ibrix_exportfs` command to create an access entry:

```
<installdirectory>/bin/ibrix_exportfs -c -f FSNAME -p CLIENT:/PATHNAME,CLIENT2:/PATHNAME,...
```

To see all access entries that have been created, use the following command:

```
<installdirectory>/bin/ibrix_exportfs -c -l
```

To remove an access entry, use the following command:

```
<installdirectory>/bin/ibrix_exportfs -c –U -f FSNAME -p CLIENT:/PATHNAME,
CLIENT2:/PATHNAME,...
```

Using Export Control

When Export Control is enabled on a file system, by default, X9000 clients have no access to the file system. Instead, the system administrator grants the clients access by executing the `ibrix_mount` command on the management console.

Enabling Export Control does not affect access from a file serving node to a file system (and thereby, NFS/CIFS client access). File serving nodes always have RW access.

To determine whether Export Control is enabled, run `ibrix_fs -i` or `ibrix_fs -l`. The output indicates whether Export Control is enabled.

To enable Export Control, include the `-C` option in the `ibrix_fs` command:

```
<installdirectory>/bin/ibrix_fs -C -E -f FSNAME
```

To disable Export Control, execute the `ibrix_fs` command with the `-C` and `-D` options:

```
<installdirectory>/bin/ibrix_fs -C -D -f FSNAME
```

To mount a file system that has Export Control enabled, include the `ibrix_mount -o {RW|RO}` option to specify that all clients have either RO or RW access to the file system. The default is RO. In addition, when specifying a hostgroup, the root user can be limited to RO access by adding the `root_ro` parameter.

3 Setting up quotas

Quotas can be assigned to individual users or groups, or to a directory tree. Individual quotas limit the amount of storage or the number of files that a user or group can use in a file system. Directory tree quotas limit the amount of storage and the number of files that can be created on a file system located at a specific directory tree.

Although it is best to set up quotas when you create a file system, you can configure them at any time. You can assign quotas to a user, group, or directory on the management console or from the CLI. You can also import quota information from a file.

If a user has a user quota and a group quota for the same file system, the first quota reached takes precedence.

The existing quota configuration can be exported to a file at any time.

NOTE: HP recommends that you export the quota configuration and save the resulting file whenever you update quotas on your cluster.

How quotas work

A quota is delimited by hard and soft storage limits for both the megabytes of storage and the number of files allotted to a user, group, or directory tree. The hard limit specifies the maximum allotted storage in terms of file size and number of files. The soft limit specifies the number of megabytes or files that, when reached, causes the file serving node to start a countdown timer.

The timer runs until either the hard storage limit is reached or seven days elapse. When the timer stops, the user, group, or directory tree for which the quota was set cannot store any more data, and the system issues `Disk quota exceeded` messages at each write attempt.

NOTE: Quota statistics are updated on a regular basis (at one-minute intervals). At each update, the file and storage usage for each quota-enabled user, group, or directory tree is queried, and the result is distributed to all file serving nodes. *Users or groups can temporarily exceed their quota if the allocation policy in effect for a file system causes their data to be written to different file serving nodes during the statistics update interval.* In this situation, it is possible for the storage usage visible to each file serving node to be below or at the quota limit while the aggregate storage use exceeds the limit.

There is a delay of several minutes between the time a command to update quotas is executed and when the results are displayed by the `ibrix_edquota -l` command. This is normal behavior.

Enabling quotas on a file system

Before you can set quota limits, quotas must be enabled. If you did not enable quotas when you created the file system, unmount the file system and then take one of these actions:

- On the management console, select the file system and then select **Quotas** from the lower Navigator. On the Quota Summary page, click **Enable**.
- From the CLI, run the following command:

  `<installdirectory>/bin/ibrix_fs -q -E -f FSNAME`

Setting user and group quotas

Before configuring quotas, the quota feature must be enabled on the file system and the file system must be mounted.

NOTE: For the purpose of setting quotas, no UID or GID can exceed 2,147,483,647.

Setting user quotas to zero removes the quotas.

NOTE: When a new NIS user is added, you need to restart the Fusionmanager services before assigning quotas to the user:

```
/etc/init.d/ibrix_fusionmanager restart
```

GUI procedure

To configure a user quota, select the file system where the quotas will be configured. Next, select **Quotas > User Quotas** from the lower Navigator, and then, on the User Quota Usage Limits page, click **Set**. User quotas can be specified by either the user name or ID. Specifying quota limits is optional.

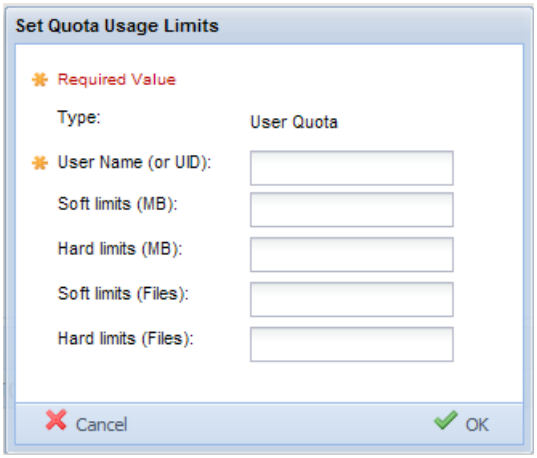

To configure a group quota, select the file system where the quotas will be configured. Next, select **Quotas > Group Quotas** from the lower Navigator, and then, on the Group Quota Usage Limits page, click **Set**. Group quotas can be identified by either the group name or GID. Specifying quota limits is optional.

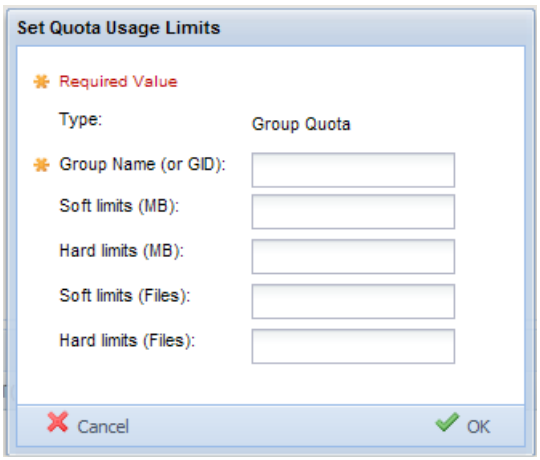

To change user or group quotas, select the appropriate user or group on the Quota Usage Limits page, and then click **Modify**.

CLI procedure

Use the following commands to set quotas for users and groups:

- Set a quota for a single user:

```
<installdirectory>/bin/ibrix_edquota -s -u "USER" -f FSNAME
[-M SOFT_MEGABYTES] [-m HARD_MEGABYTES] [-I SOFT_FILES] [-i HARD_FILES]
```

- Set a quota for a single group:

```
<installdirectory>/bin/ibrix_edquota -s -g "GROUP" -f FSNAME [-M SOFT_MEGABYTES]
[-m HARD_MEGABYTES] [-I SOFT_FILES] [-i HARD_FILES]
```

Enclose the user or group name in single or double quotation marks.

Setting directory tree quotas

Directory tree quotas limit the amount of storage and the number of files that can be created on a file system located at a specific directory tree. Before configuring directory tree quotas, the quota feature must be enabled on the file system and the file system must be mounted.

GUI procedure

To configure a directory tree quota, select the file system where the quotas will be configured. Next, select **Quotas > Directory Quotas** from the lower Navigator, and then, on the Directory Tree Quota Usage Limits page, click **Create**. For **Name (Alias)**, enter a unique name for the directory tree quota. The name cannot contain a comma (,) character.

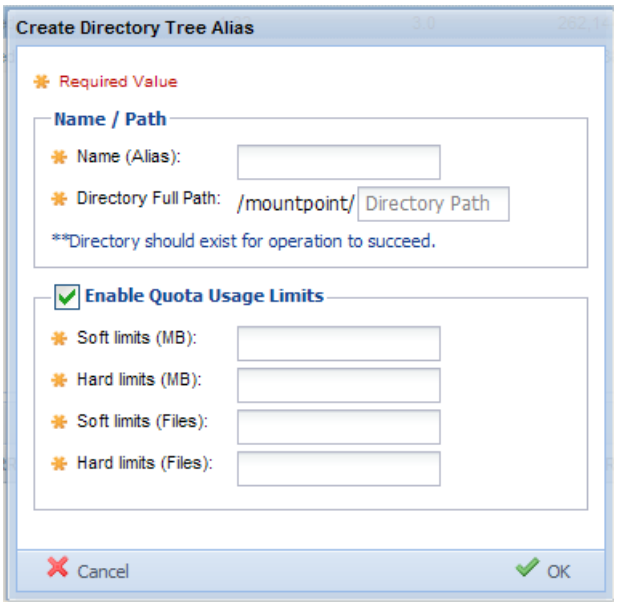

To change a directory tree quota, select the directory tree on the Quota Usage Limits page, and then click **Modify**.

CLI procedure

To assign quotas to a directory tree, complete these steps:

1. Create the directory tree quota:

   ```
   ibrix_fs_ops -D -c -f FSNAME -p PATH -n NAME
   ```

   The `-f FSNAME` option specifies the name of the file system. The `-p PATH` option specifies the pathname of the directory tree. If the pathname includes a space, enclose the portion of the pathname that includes the space in single quotation marks, and enclose the entire pathname in double quotation marks. For example:

   ```
   ibrix_fs_ops -D -c -f fs48 -p "/fs48/data/'QUOTA 4'" -n QUOTA_4
   ```

The `-n NAME` option specifies a unique name for the directory tree quota. The name cannot contain a comma (,) character.

2. Assign usage limits to the directory tree quota:

   ```
   ibrix_edquota -s -d NAME -f FSNAME -M SOFT_MEGABYTES -m
   HARD_MEGABYTES -I SOFT_FILES -i HARD_FILES
   ```

   The `-d NAME` option specifies the name of the directory tree quota. The `-f FSNAME` option specifies the name of the file system. Use `-M SOFT_MEGABYTES` and `-m HARD_MEGABYTES` to specify soft and hard limits for the megabytes of storage allowed on the directory tree. Use `-I SOFT_FILES` and `-i HARD_FILES` to specify soft and hard limits for the number of files allowed on the directory tree.

Using a quotas file

Quota limits can be imported into the cluster from the quotas file, and existing quotas can be exported to the file. See "Format of the quotas file" (page 22) for the format of the file.

Importing quotas from a file

From the management console, select the file system, select **Quotas** from the lower Navigator, and then click **Import**.

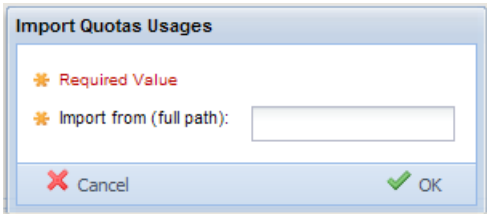

From the CLI, use the following command to import quotas from a file, where *PATH* is the path to the quotas file:

```
ibrix_edquota -t -p PATH [-f FSNAME]
```

Exporting quotas to a file

From the management console, select the file system, select **Quotas** from the lower Navigator, and then click **Export**.

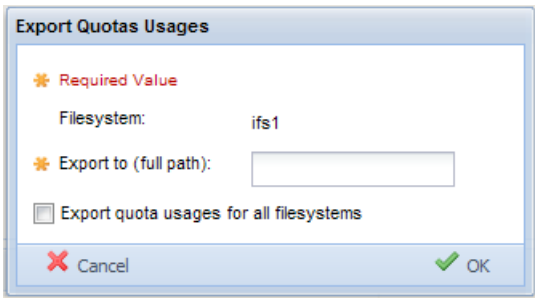

From the CLI, use the following command to export the existing quotas information to a file, where *PATH* is the pathname of the quotas file:

```
ibrix_edquota -e -p PATH [-f FSNAME]
```

Format of the quotas file

The quotas file contains a line for each user, group, or directory tree assigned a quota. The lines must use one of the following formats. The "A" format specifies a user or group ID. The "B" format

specifies a user or group name, or a directory tree that has already been assigned an identifier name. The "C" format specifies a directory tree, where the path exists, but the identifier name for the directory tree will not be created until the quotas are imported.

```
A,{type},{block_hardlimit},{block_soft-limit},{inode_hardlimit},{inode_softlimit},{id}
B,{type},{block_hardlimit},{block_soft-limit},{inode_hardlimit},{inode_softlimit},"{name}"
C,{type},{block_hardlimit},{block_soft-limit},{inode_hardlimit},{inode_softlimit},
"{name}","{path}"
```

The fields in each line are:

`{type}`
: Either `0` for a user quota; `1` for a group quota; 2 for a directory tree quota.

`{block_hardlimit}`
: The maximum number of 1K blocks allowed for the user, group, or directory tree. (1 MB = 1024 blocks).

`{block_soft-limit}`
: The number of 1K blocks that, when reached, starts the countdown timer.

`{inode_hardlimit}`
: The maximum number of files allowed for the user, group, or directory tree.

`{inode_softlimit}`
: The number of files that, when reached, starts the countdown timer.

`{id}`
: The UID for a user quota or the GID for a group quota.

`{name}`
: A user name, group name, or directory tree identifier.

`{path}`
: The full path to the directory tree. The path must already exist.

Using online quota check

You will need to rescan the quota usage for a user, group, or directory tree in the following cases:

- You turned quotas off for a user, the user continued to store data in a file system, and you now want to turn quotas back on for this user.
- You are setting up quotas for the first time for a user who has previously stored data in a file system.
- You are using directory tree quotas and you have moved a subdirectory into another parent directory outside of the directory that has the directory tree quota.

The `ibrix_online_quotacheck` command is used to rescan quota usage. The command must be run from a file serving node that has the file system mounted.

To run a quota check on a file system, use the following command:

```
ibrix_online_quotacheck –M {quotamonitor_host} /mountpoint
```

For the `–M` option, it is best to specify the management console as the target host/IP address. This allows the quota usage for a deleted user or group to be updated. You can specify pathnames in place of *`/mountpoint`* if you want to target specific directories instead of an entire file system.

To use the `–M` option for a directory tree quota, include the `-P` option in the command:

```
ibrix_online_quotacheck –M {quotamonitor_host} -P /mountpoint
```

To run a quota check on a specific user having files only in the home directory, use the following command:

```
ibrix_online_quotacheck -u –T userid home_dir
```

To run a quota check on a specific group, use the following command:

```
ibrix_online_quotacheck -g –T groupid home_dir
```

To run a quota check on a specific directory tree, use the following command:

```
ibrix_online_quotacheck -t –P dirtree_path
```

To remove the directory tree quota for all files under a specific path, use the following command:

```
ibrix_online_quotacheck -t –T 0 dirtree_path
```

Configuring email notifications for quota events

If you would like to be notified when certain quota events occur, you can set up email notification for those events. On the Management Console GUI, select **Email Configuration**. In the Events Notified by Email tab, select the appropriate events and specify the email addresses to be notified.

Deleting quotas

Quotas can be deleted from the management console or the CLI. From the management console, select the quota from the appropriate Quota Usage Limits page and then click **Delete**. To delete quotas from the CLI, use the following procedures.

Deleting user and group quotas

You can delete user and group quotas at any time.

To delete quotas for a user, use the following command:

```
<installdirectory>/bin/ibrix_edquota -D -u UID [-f FSNAME]
```

To delete quotas for a group, use the following command:

```
<installdirectory>/bin/ibrix_edquota -D -g GID [-f FSNAME]
```

Deleting directory tree quotas or usage limits

Use the following commands to delete directory tree quotas or their usage limits.

- Delete usage limits for a directory tree quota:

  ```
  ibrix_edquota -D -d NAME -f FSNAME
  ```

  The `-d NAME` option specifies the name of the directory tree quota.

- Delete a directory tree quota for a specific file system:

  1. Issue the following command to delete the directory tree quota entry and limits:

     ```
     ibrix_fs_ops -D -d -f FSNAME -n NAME
     ```

     The `-n NAME` option specifies the name of the directory tree quota.

  2. Issue the following command to remove the quota account for the directory tree:

     ```
     ibrix_online_quotacheck -t -T 0 {path}
     ```

4 Maintaining file systems

This chapter describes how to extend a file system, rebalance segments, delete a file system or file system component, and check or repair a file system. The chapter also includes file system troubleshooting information.

Viewing information about file systems and components

The Filesystems page on the management console GUI displays comprehensive information about a file system and its components. This section describes how to view the same information from the command line.

Viewing physical volume information

Use the following command to view information about physical volumes:

```
<installdirectory>/bin/ibrix_pv -l
```

The following table lists the output fields for `ibrix_pv -l`.

Field	Description
PV_Name	Physical volume name. Regular physical volume names begin with the letter *d*. The names of physical volumes that are part of a mirror device begin with the letter *m*. Both are numbered sequentially.
Size (MB)	Physical volume size, in MB.
VG name	Name of volume group created on this physical volume, if any.
RAID type	Not applicable for this release.
RAID host	Not applicable for this release.
RAID device	Not applicable for this release.
Network host	Not applicable for this release.
Network port	Not applicable for this release.

Viewing volume group information

To display summary information about all volume groups, use the `ibrix_vg -l` command:

```
<installdirectory>/bin/ibrix_vg -l
```

The VG_FREE field indicates the amount of group space that is not allocated to any logical volume. The VG_USED field reports the percentage of available space that is allocated to a logical volume.

To display detailed information about volume groups, use the `ibrix_vg -i` command. The `-g VGLIST` option restricts the output to the specified volume groups.

`<installdirectory>/bin/ibrix_vg -i [-g VGLIST]`

The following table lists the output fields for `ibrix_vg -i`.

Field	Description
Name	Volume group name.
Size (MB)	Volume group size in MB.
Free (MB)	Free (unallocated) space, in MB, available on this volume group.
Used (percentage)	Percentage of total space in the volume group allocated to logical volumes.
File System Name	File system to which this logical volume belongs.
Physical Volume Name	Name of the physical volume used to create this volume group.
Physical Volume Size	Size, in MB, of the physical volume used to create this volume group.
Logical Volume Name	Names of logical volumes created from this volume group.
Logical Volume Size	Size, in MB, of each logical volume created from this volume group.
File System Generation	Number of times the structure of the file system has changed (for example, new segments were added).
Segment Number	Number of this segment (logical volume) in the file system.
Host Name	File serving node that owns this logical volume.
State	Operational state of the file serving node. See the administration guide for your system for a list of the states.

Viewing logical volume information

To view information about logical volumes, use the `ibrix_lv -l` command:

`<installdirectory>/bin/ibrix_lv -l`

The following table lists the output fields for `ibrix_lv -l`.

Field	Description
LV_NAME	Logical volume name.
LV_SIZE	Logical volume size, in MB.
FS_NAME	File system to which this logical volume belongs.
SEG_NUM	Number of this segment (logical volume) in the file system.
VG_NAME	Name of the volume group created on this physical volume, if any.
OPTIONS	Linux `lvcreate` options that have been set on the volume group.

Viewing file system information

To view information about all file systems, use the `ibrix_fs -l` command. This command also displays information about any file system snapshots.

`<installdirectory>/bin/ibrix_fs -l`

The following table lists the output fields for `ibrix_fs -l`.

Field	Description
FS_NAME	File system name.
STATE	State of the file system (for example, Mounted).
CAPACITY (GB)	Total space available in the file system, in GB.
USED%	Amount of space used in the file system.
Files	Number of files that can be created in this file system.
FilesUsed%	Percentage of total storage used by files and directories.
GEN	Number of times the structure of the file system has changed (for example, new segments were added).
NUM_SEGS	Number of file system segments.

To view detailed information about file systems, use the `ibrix_fs -i` command. To view information for all file systems, omit the `-f FSLIST` argument.

`<installdirectory>/bin/ibrix_fs -i [-f FSLIST]`

The following table lists the file system output fields reported by `ibrix_fs -i`.

Field	Description
Total Segments	Number of segments.
STATE	State of the file system (for example, Mounted).
Mirrored?	Not applicable for this release.
Compatible?	`Yes` indicates that the file system is 32-bit compatible; the maximum number of segments (maxsegs) allowed in the file system is also specified. `No` indicates a 64-bit file system.
Generation	Number of times the structure of the file system has changed (for example, new segments were added).
FS_ID	File system ID for NFS access.
FS_NUM	Unique X9000 Software internal file system identifier.
EXPORT_CONTROL_ENABLED	`Yes` if enabled; `No` if not.
QUOTA_ENABLED	`Yes` if enabled; `No` if not.
DEFAULT_BLOCKSIZE	Default block size, in KB.
CAPACITY	Capacity of the file system.
FREE	Amount of free space on the file system.
AVAIL	Space available for user files.
USED PERCENT	Percentage of total storage occupied by user files.
FILES	Number of files that can be created in this file system.
FFREE	Number of unused file inodes available in this file system.
Prealloc	Number of KB a file system preallocates to a file; default: 1,024 KB.
Readahead	Number of KB that X9000 Software will pre-fetch; default: 512 KB.
NFS Readahead	Number of KB that X9000 Software pre-fetches under NFS; default: 256 KB.

Field	Description
Default policy	Allocation policy assigned on this file system. Defined policies are: ROUNDROBIN, STICKY, DIRECTORY, LOCAL, RANDOM, and NONE. See "File allocation policies" (page 116) for information on these policies.
Default start segment	The first segment to which an allocation policy is applied in a file system. If a segment is not specified, allocation starts on the segment with the most storage space available.
File replicas	NA.
Dir replicas	NA.
Mount Options	Possible root segment inodes. This value is used internally.
Root Segment Hint	Current root segment number, if known. This value is used internally.
Root Segment Replica(s) Hint	Possible segment numbers for root segment replicas. This value is used internally.
Snap FileSystem Policy	Snapshot strategy, if defined.

The following table lists the per-segment output fields reported by `ibrix_fs -i`.

Field	Description
SEGMENT	Number of segments.
OWNER	The host that owns the segment.
LV_NAME	Logical volume name.
STATE	The current state of the segment (for example, `OK` or `UsageStale`).
BLOCK_SIZE	Default block size, in KB.
CAPACITY (GB)	Size of the segment, in GB.
FREE (GB)	Free space on this segment, in GB.
AVAIL (GB)	Space available for user files, in GB.
FILES	Inodes available on this segment.
FFREE	Free inodes available on this segment.
USED%	Percentage of total storage occupied by user files.
BACKUP	Backup host name.
TYPE	Segment type. `MIXED` means the segment can contain both files and directories.
TIER	Tier to which the segment was assigned.
LAST_REPORTED	Last time the segment state was reported.
HOST_NAME	Host on which the file system is mounted.
MOUNTPOINT	Host mountpoint.
PERMISSION	File system access privileges: RO or RW.
Root_RO	Specifies whether the root user is limited to read-only access, regardless of the access setting.

Lost+found directory

When browsing the contents of X9000 Software file systems, you will see a directory named `lost+found`. This directory is required for file system integrity and should not be deleted.

Viewing disk space information from a Linux X9000 client

Because file systems are distributed among segments on many file serving nodes, disk space utilities such as `df` must be provided with collated disk space information about those nodes. The management console collects this information periodically and collates it for `df`.

X9000 Software includes a disk space utility, `ibrix_df`, that enables Linux X9000 clients to obtain utilization data for a file system. Execute the following command on any Linux X9000 client:

`<installdirectory>/bin/ibrix_df`

The following table lists the output fields for `ibrix_df`.

Field	Description
Name	File system name.
CAPACITY	Number of blocks in the file system.
FREE	Number of unused blocks of storage.
AVAIL	Number of blocks available for user files.
USED PERCENT	Percentage of total storage occupied by user files.
FILES	Number of files that can be created in the file system.
FFREE	Number of unused file inodes in the file system.

Extending a file system

You can extend a file system from the Management Console GUI or the CLI.

Select the file system on the Filesystems page and click **Extend**. The Extend Filesystem dialog box allows you to select the storage to be added to the file system. If data tiering is used on the file system, you can also enter the name of the appropriate tier.

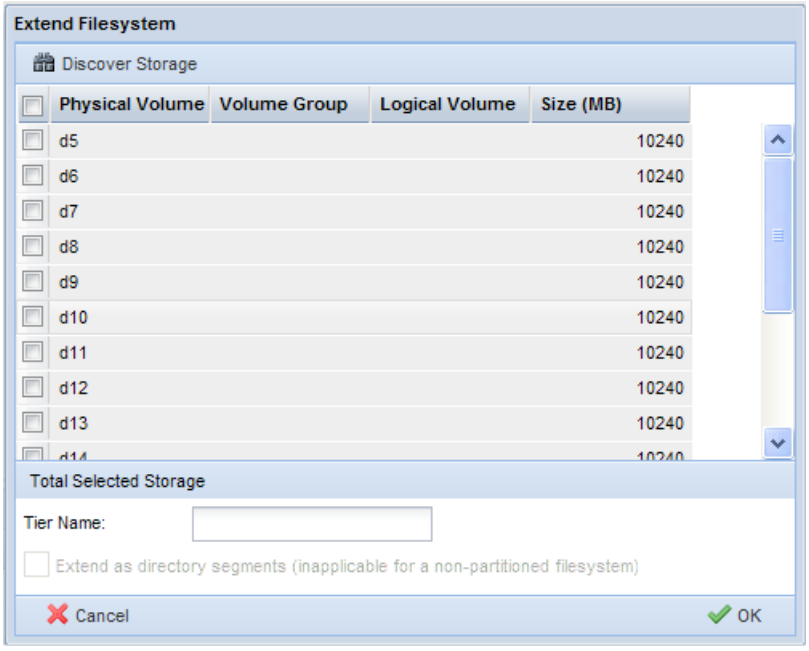

On the CLI, use the `ibrix_fs` command to extend a file system. Segments are added to the file serving nodes in a round-robin manner. If tiering rules are defined for the file system, the `-t` option is required. Avoid expanding a file system while a tiering job is running. The expansion takes priority and the tiering job is terminated.

Extend a file system with the logical volumes (segments) specified in *`LVLIST`*:

```
<installdirectory>/bin/ibrix_fs -e -f FSNAME -s LVLIST [-t TIERNAME]
```

Extend a file system with segments created from the physical volumes in *LVLIST*:

```
<installdirectory>/bin/ibrix_fs -e -f FSNAME -p PVLIST [-t TIERNAME]
```

Extend a file system with specific logical volumes on specific file serving nodes:

```
<installdirectory>/bin/ibrix_fs -e -f FSNAME -S LV1:HOSTNAME1,LV2:HOSTNAME2...
```

Rebalancing segments in a file system

Segment rebalancing involves redistributing files among segments in a file system to balance segment utilization and server workload. For example, after adding new segments to a file system, you can rebalance all segments to redistribute files evenly among the segments. Usually, you will want to rebalance all segments, possibly as a `cron` job. In special situations, you might want to rebalance specific segments. Segments marked as bad (that is, segments that cannot be activated for some reason) are not candidates for rebalancing.

The rebalancing feature can also evacuate segments (or logical volumes) located on storage that will be removed from the cluster. This procedure moves the data to other segments in the file system, and is transparent to users or applications accessing the file system. For more information about removing storage from a cluster, see the administration guide for your product.

A file system must be mounted when you rebalance its segments.

How rebalancing works

During a rebalance operation on a file system, files are moved from source segments to destination segments. X9000 Software calculates the average aggregate utilization of the selected source segments, and then moves files from sources to destinations to bring each candidate source segment as close as possible to the calculated utilization threshold. The final absolute percent usage in the segments depends on the average file size for the target file system. If you do not specify any sources or destinations for a rebalance task, candidate segments are sorted into sources and destinations and then rebalanced as evenly as possible.

If you specify sources, all other candidate segments in the file system are tagged as destinations, and vice versa if you specify destinations. Following the general rule, X9000 Software will calculate the utilization threshold from the sources, and then bring the sources as close as possible to this value by evenly distributing their excess files among all destinations. If you specify sources, only those segments are rebalanced, and the overflow is distributed among all remaining candidate segments. If you specify destinations, all segments except the specified destinations are rebalanced, and the overflow is distributed only to the destinations. If you specify both sources and destinations, only the specified sources are rebalanced, and the overflow is distributed only among the specified destinations.

If there is not enough aggregate room in destination segments to hold the files that must be moved from source segments in order to balance the sources, X9000 Software issues an error message and does not move any files. The more restricted the number of destinations, the higher the likelihood of this error.

When rebalancing segments, note the following:

- To move files out of certain overused segments, specify source segments.
- To move files into certain underused segments, specify destination segments.
- To move files out of certain segments and place them in certain destinations, specify both source and destination segments.

Rebalancing segments on the management console GUI

Select the file system on the management console GUI and then select **Tasks** > **Rebalancer** from the lower Navigator. Click **Start** on the Task Summary screen to open the Start Rebalancing dialog. The General tab allows you to select analytical mode.

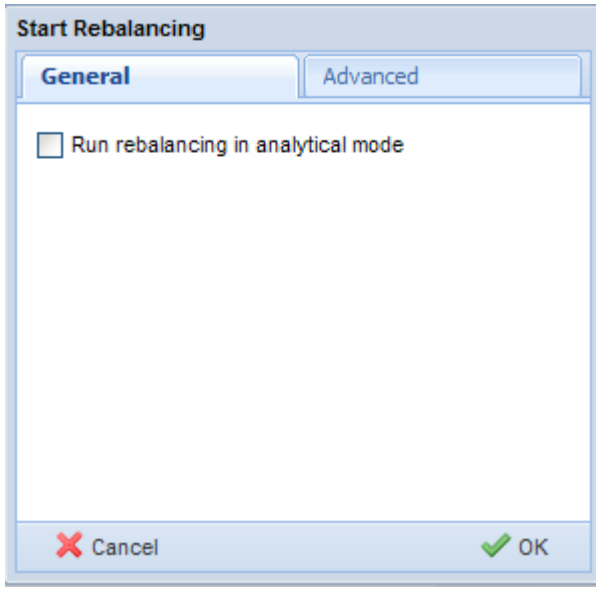

On the Advanced tab, select the source segments, destination segments, or both for the rebalancing task. You can also specify segments to evacuate.

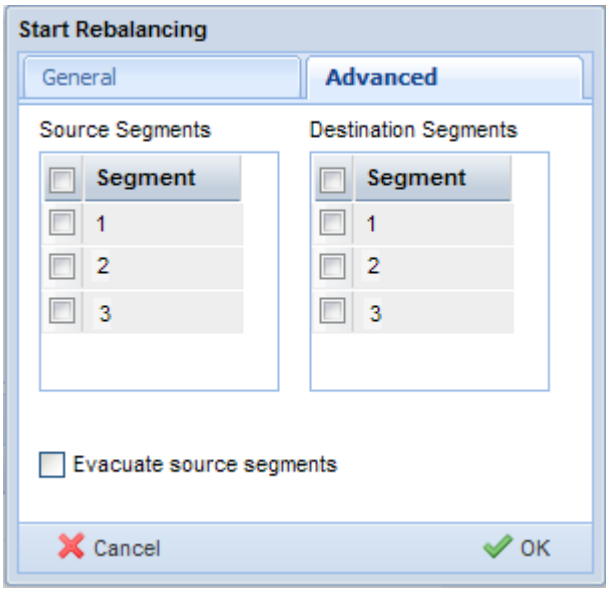

Rebalancing segments from the CLI

To rebalance all segments, use the following command. Include the `-a` option to run the rebalance operation in analytical mode.

```
<installdirectory>/bin/ibrix_rebalance -r -f FSNAME
```

To rebalance by specifying specific source segments, use the following command:

```
<installdirectory>/bin/ibrix_rebalance -r -f FSNAME [[-s SRCSEGMENTLIST] [-S SRCLVLIST]]
```

For example, to rebalance segments 2 and 3 only and to specify them by segment name:

```
<installdirectory>/bin/ibrix_rebalance -r -f ifs1 -s 2,3
```

To rebalance segments 1 and 2 only and to specify them by their logical volume names:

```
<installdirectory>/bin/ibrix_rebalance -r -f ifs1 -S ilv1,ilv2
```

To rebalance by specifying specific destination segments, use the following command:

```
<installdirectory>/bin/ibrix_rebalance -r -f FSNAME [[-d DESTSEGMENTLIST] [-D DESTLVLIST]]
```

For example, to rebalance segments 3 and 4 only and to specify them by segment name:

```
<installdirectory>/bin/ibrix_rebalance -r -f ifs1 -d 3,4
```

To rebalance segments 3 and 4 only and to specify them by their logical volume names:

```
<installdirectory>/bin/ibrix_rebalance -r -f ifs1 -D ilv3,ilv4
```

Tracking the progress of a rebalance task

You can use the management console GUI or CLI to track the progress of a rebalance task. As a rebalance task progresses, usage approaches an average value across segments, excluding bad segments that are not candidates for rebalancing or segments containing files that are in heavy use during the operation.

To track the progress of a rebalance task on the management console GUI, select the file system, and then select **Rebalancer** from the lower Navigator. The Task Summary displays details about the rebalance task. Also examine `Used (%)` on the Segments screen for the file system.

To track rebalance job progress from the CLI, use the `ibrix_fs -i` command:

```
<installdirectory>/bin/ibrix_fs -i
```

The output lists detailed information about the file system. The `USED%` field shows usage per segments.

Viewing the status of rebalance tasks

Use the following commands to view status for jobs on all file systems or only on the file systems specified in *FSLIST*:

```
<installdirectory>/bin/ibrix_rebalance -l [-f FSLIST]

<installdirectory>/bin/ibrix_rebalance -i [-f FSLIST]
```

The first command reports summary information. The second command lists jobs by task ID and file system and indicates whether the job is running or stopped. Jobs that are in the analysis (Coordinator) phase are listed separately from those in the implementation (Worker) phase.

Stopping rebalance tasks

You can stop running or stalled rebalance tasks. If the management console cannot stop the task for some reason, you can force the task to stop. Stopping a task poses no risks for the file system. The management console completes any file migrations that are in process when you issue the stop command. Depending on when you stop a task, segments might contain more or fewer files than before the operation started.

To stop a rebalance task on the management console GUI, select the file system, and then select **Rebalancer** from the lower Navigator. Click **Stop** on the Task Summary to stop the task.

To stop a task from the CLI, first execute `ibrix_rebalance -i` to obtain the *TASKID*, and then execute the following command:

```
<installdirectory>/bin/ibrix_rebalance -k -t TASKID [-F]
```

To force the task to stop, include the `-F` option.

Disabling 32-bit mode on a file system

If your cluster clients are converting from 32-bit to 64-bit applications, you can disable 32-bit mode on the file system, which enables 64-bit mode. (For information about 64-bit mode, see "Using 32-bit or 64-bit mode" (page 11).)

To determine whether 64-bit mode is enabled on a file system, execute the command `ibrix_fs -i`. If the output reports `Compatible? : No`, 64-bit mode is enabled.

NOTE: A file system using 64-bit mode cannot be changed to use 32-bit mode. If there is a chance that clients will need to run a 32-bit application, do not disable 32-bit mode.

To disable 32-bit mode, complete these steps:

1. Unmount the file system.
2. On the GUI, select the file system and click **Modify** on the Summary tab. On the Modify Filesystems Properties dialog box, select **Disable 32 Bit Compatibility Mode**.

   From the CLI, execute the following command:

   ```
   <installdirectory>/bin/ibrix_fs -w -f FSNAME
   ```
3. Remount the file system.

Deleting file systems and file system components

Deleting a file system

Before deleting a file system, unmount it from all file serving nodes, X9000 clients, NFS clients, and CIFS clients. (See "Unmounting a file system" (page 16).)

CAUTION: When a file system is deleted from the configuration database, its data becomes inaccessible. To avoid unintended service interruptions, be sure you have specified the correct file system.

To delete a file system, use the following command:

```
<installdirectory>/bin/ibrix_fs -d -f FSLIST
```

For example, to delete file systems `ifs1` and `ifs2`:

```
<installdirectory>/bin/ibrix_fs -d -f ifs1,ifs2
```

Deleting segments, volume groups, and physical volumes

When deleting segments, volume groups, or physical volumes, you should be aware of the following:

- A segment cannot be deleted until the file system to which it belongs is deleted.
- A volume group cannot be deleted until all segments that were created on it are deleted.
- A physical volume cannot be deleted until all volume groups created on it are deleted.

If you delete physical volumes but do not remove the physical storage from the network, the volumes might be rediscovered when you next perform a discovery scan on the cluster.

To delete segments:

```
<installdirectory>/bin/ibrix_lv -d -s LVLIST
```

For example, to delete segments `ilv1` and `ilv2`:

```
<installdirectory>/bin/ibrix_lv -d -s ilv1,ilv2
```

To delete volume groups:

```
<installdirectory>/bin/ibrix_vg -d -g VGLIST
```

For example, to delete volume groups `ivg1` and `ivg2`:

```
<installdirectory>/bin/ibrix_vg -d -g ivg1,ivg2
```

To delete physical volumes:

```
<installdirectory>/bin/ibrix_pv -d -p PVLIST [-h HOSTLIST]
```

For example, to delete physical volumes `d1`, `d2`, and `d3`:

```
<installdirectory>/bin/ibrix_pv -d -p d[1-3]
```

Deleting file serving nodes and X9000 clients

Before deleting a file serving node, unmount all file systems from it and migrate any segments that it owns to a different server. Ensure that the file serving node is not serving as a failover standby and is not involved in network interface monitoring. To delete a file serving node, use the following command:

```
<installdirectory>/bin/ibrix_server -d -h HOSTLIST
```

For example, to delete file serving nodes s1.hp.com and s2.hp.com:

```
<installdirectory>/bin/ibrix_server -d -h s1.hp.com,s2.hp.com
```

To delete X9000 clients, use the following command:

```
<installdirectory>/bin/ibrix_client -d -h HOSTLIST
```

Checking and repairing file systems

CAUTION: Do not run `ibrix_fsck` in corrective mode without the direct guidance of HP Support. If run improperly, the command can cause data loss and file system damage.

CAUTION: Do not run `e2fsck` (or any other off-the-shelf `fsck` program) on any part of a file system. Doing this can damage the file system.

The `ibrix_fsck` command can detect and repair file system inconsistencies, which are a symptom of file system corruption. File system inconsistencies can occur for many reasons, including hardware failure, power failure, switching off the system without proper shutdown, and failed migration.

The command runs in four phases and has two running modes: analytical and corrective. You must run the phases in order and you must run all of them:

- Phase 0 checks host connectivity and the consistency of segment byte blocks and repairs them in corrective mode.
- Phase 1 checks segments and repairs them in corrective mode. Results are stored locally.
- Phase 2 checks the file system and repairs it in corrective mode. Results are stored locally.
- Phase 3 moves files from `lost+found` on each segment to the global `lost+found` directory on the root segment for the file system.

If a file system shows evidence of corruption, contact HP Support. A representative will ask you to run `ibrix_fsck` in analytical mode and, based on the output, will recommend a course of action and assist in running the command in corrective mode. HP strongly recommends that you use corrective mode only with the direct guidance of HP Support. Corrective mode is complex and difficult to run safely. Using it improperly can damage both data and the file system. Analytical mode is completely safe, by contrast.

NOTE: During an `ibrix_fsck` job, an INFSCK flag is set on the file system to protect it. If an error occurs during the job, you must explicitly clear the INFSCK flag (see "Clearing the INFSCK flag on a file system" (page 35)), or you will be unable to mount the file system.

Analyzing the integrity of a file system on all segments

Use the following procedure to analyze file system integrity:

1. Turn off automated failover (see "Turning automated failover on and off" in the administration guide for your system), unmount all NFS clients and stop NFS, and then unmount the file system.

2. Execute the following commands one time each:

```
<installdirectory>/bin/ibrix_fsck -f FSNAME -p 0 [-s LVNAME]

<installdirectory>/bin/ibrix_fsck -f FSNAME -p 1 [-s LVNAME] [-B BLOCKSIZE]
[-b ALTSUPERBLOCK]

<installdirectory>/bin/ibrix_fsck -f FSNAME -p 2 -m MOUNTPOINT

<installdirectory>/bin/ibrix_fsck -f FSNAME -p 3 -m MOUNTPOINT
```

Clearing the INFSCK flag on a file system

To clear the INFSCK flag, use the following command:

```
<installdirectory>/bin/ibrix_fsck -f FSNAME -C
```

Troubleshooting file systems

ibrix_pv -a discovers too many or too few devices

This situation occurs when file serving nodes see devices multiple times. To prevent this, modify the LVM2 filter in `/etc/lvm/lvm.conf` to filter only on devices used by X9000 Software. This will change the output of `lvmdiskscan`.

By default, the following filter finds all devices:

```
filter = [ "a/.*/" ]
```

The following filter finds all sd devices:

```
filter = [ "a|^/dev/sd.*|", "r|^.*|" ]
```

Contact HP Support if you need assistance.

Cannot mount on an X9000 client

Verify the following:

- The file system is mounted and functioning on the file serving nodes.
- The mountpoint exists on the X9000 client. If not, create the mountpoint locally on the client.
- Software management services have been started on the X9000 client (see "Starting and stopping processes" in the administration guide for your platform).

NFS clients cannot access an exported file system

An exported file system has been unmounted from one or more file serving nodes, causing X9000 Software to automatically disable NFS on those servers. Fix the issue causing the unmount and then remount the file system.

User quota usage data is not being updated

Restart the quota monitor service to force a read of all quota usage data and update usage counts to the file serving nodes in your cluster. Use the following command:

```
<installdirectory>/bin/ibrix_qm restart
```

SegmentNotAvailable is reported

When writes to a segment do not succeed, the segment status may change to `SegmentNotAvailable` on the GUI and an alert message may be generated. To correct this situation, take the following steps:

1. Identify the file serving node that owns the segment. This information is reported on the Filesystem Segments page on the GUI.
2. Fail over the file serving node to its standby. See the administration guide for your system for more information about this procedure.
3. Reboot the file serving node.
4. When the file serving node is up, verify that the segment, or LUN, is available.

If the segment is still not available, contact HP Support.

SegmentRejected is reported

This alert is generated by a client call for a segment that is no longer accessible by the segment owner or file serving node specified in the client's segment map. The alert is logged to the `Iad.log` and messages files. It is usually an indication of an out-of-date or stale segment map for the affected file system and is caused by a network condition. Other possible causes are rebooting the node, unmounting the file system on the node, segment migrations, and, in a failover scenario, stale IAD, an unresponsive kernel, or a network RPC condition.

To troubleshoot this alert, check network connectivity among the nodes, ensuring that the network is optimal and any recent network conditions have been resolved. From the file system perspective, verify segment maps by comparing the file system generation numbers and the ownership for those segments being rejected by the clients.

Use the following commands to compare the file system generation number on the local file serving nodes and the clients logging the error.

`/usr/local/ibrix/bin/rtool enumseg <FSNAME> <SEGNUMBER>`

For example:

```
rtool enumseg ibfs1 3
segnum=3 of 4 -----------
   fsid ........................... 7b3ea891-5518-4a5e-9b08-daf9f9f4c027
   fsname ......................... ibfs1
   device_name .................... /dev/ivg3/ilv3
   host_id ........................ 1e9e3a6e-74e4-4509-a843-c0abb6fec3a6
   host_name ...................... ib50-87 <-- Verify owner of segment
   ref_counter .................... 1038
   state_flags .................... SEGMENT_LOCAL SEGMENT_PREFERED SEGMENT_DHB <SEGMENT_
         ORPHAN_LIST_CREATED (0x00100061)
   write_WM ....................... 99129 4K-blocks (387 Mbytes)
   create_WM ...................... 793033 4K-blocks (3097 Mbytes)
   spillover_WM ................... 892162 4K-blocks (3485 Mbytes)
   generation ..................... 26
   quota .......................... usr,grp,dir
   f_blocks ....................... 0011895510 4K-blocks (==0047582040 1K-blocks, 46466 M)
   f_bfree ........................ 0011785098 4K-blocks (==0047140392 1K-blocks, 46035 M)
   f_bused ........................ 0000110412 4K-blocks (==0000441648 1K-blocks, 431 M)
   f_bavail ....................... 0011753237 4K-blocks (==0047012948 1K-blocks, 45911 M)
   f_files ........................ 6553600
   f_ffree ........................ 6552536
   used files (f_files - f_ffree).. 1064
          Segment statistics for 690812.89 seconds :
   n_reads=0, kb_read=0, n_writes=0, kb_written=0, n_creates=2, n_removes=0
```

Also run the following command:

`/usr/local/ibrix/bin/rtool enumfs <FSNAME>`

For example:

```
rtool enumfs ibfs1
1:----------------
   fsname ......................... ibfs1
   fsid ........................... 7b3ea891-5518-4a5e-9b08-daf9f9f4c027
   fsnum .......................... 1
   fs_flags........................ operational
   total_number_of_segments ....... 4
   mounted ........................ TRUE
   ref_counter .................... 6
   generation ..................... 26  <-- FS generation number for comparison
   alloc_policy.................... RANDOM
```

```
dir_alloc_policy.............. NONE
cur_segment................... 0
sup_ap_on..................... NONE
local_segments ............... 3
quota ........................ usr,grp,dir
f_blocks ..................... 0047582040 4K-blocks (==0190328160 1K-blocks)
f_bfree ...................... 0044000311 4K-blocks (==0176001244 1K-blocks)
f_bused ...................... 0003581729 4K-blocks (==0014326916 1K-blocks)
f_bavail ..................... 0043872867 4K-blocks (==0175491468 1K-blocks)
f_files ...................... 26214400
f_ffree ...................... 26212193
used files (f_files - f_free)... 2207
        FS statistics for 0.0 seconds :
n_reads=0, kb_read=0, n_writes=0, kb_written=0, n_creates=0, n_removes=0
```

Use the output to determine whether the FS generation number is in sync and whether the file serving nodes agree on the ownership of the rejected segments. In the `rtool enumseg` output, check the `state_flags` field for `SEGMENT_IN_MIGRATION`, which indicates that the segment is stuck in migration because of a failover.

Typically, if the segment has a healthy state flag on the file serving node that owns the segment and all file serving nodes agree on the owner of the segment, this is not a file system or file serving node issue. If a state flag is stale or indicates that a segment is in migration, call HP Support for a recovery procedure.

Otherwise, the alert indicates a file system generation mismatch. Take the following steps to resolve this situation:

1. From the active management console, run the following command to propagate a new file system segment map throughout the cluster. This step takes a few minutes.

   ```
   ibrix_dbck -I -f <FSNAME>
   ```

2. If problems persist, try restarting the client's IAD:

   ```
   /usr/local/ibrix/init/ibrix_iad restart
   ```

5 Using NFS

This section describes how to export file systems for NFS, how to autoconnect NFS clients, and how to set up NFS clients.

Exporting a file system

Exporting a file system makes local directories available for NFS clients to mount. The management console manages the table of exported file systems and distributes the information to the `/etc/exports` files on the file serving nodes. All entries are automatically re-exported to NFS clients and to the file serving node standbys unless you specify otherwise.

On the exporting file serving node, configure the number of NFS server threads based on the expected workload. The default is eight threads. If the file serving node will service many clients, you can increase the value to 16 or 64. To configure server threads, use the following command to change the default value of RPCNFSDCOUNT in the `/etc/sysconfig/nfs` file from 8 to 16 or 64.

```
ibrix_host_tune -C -h HOSTS -o nfsdCount=64
```

A file system must be mounted before it can be exported.

On the GUI

To export a file system from the GUI, select the file system, select NFS Exports from the lower Navigator panel, and then select **Export** to display the Export Filesystem via NFS dialog box.

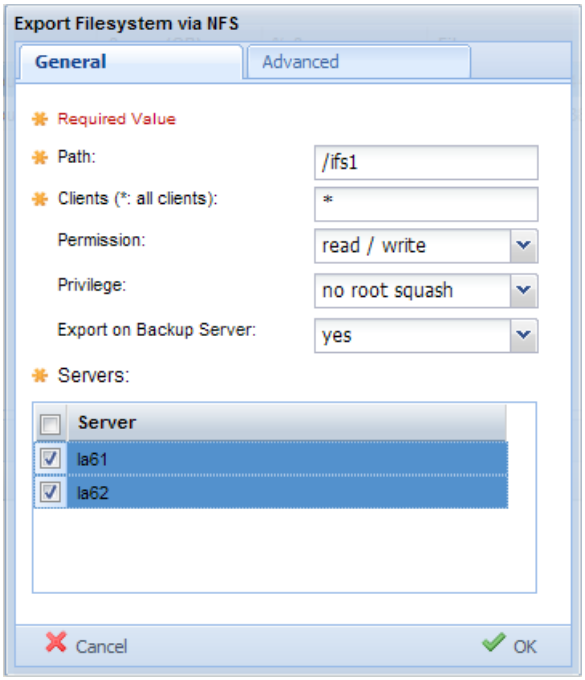

The Advanced tab displays additional options:

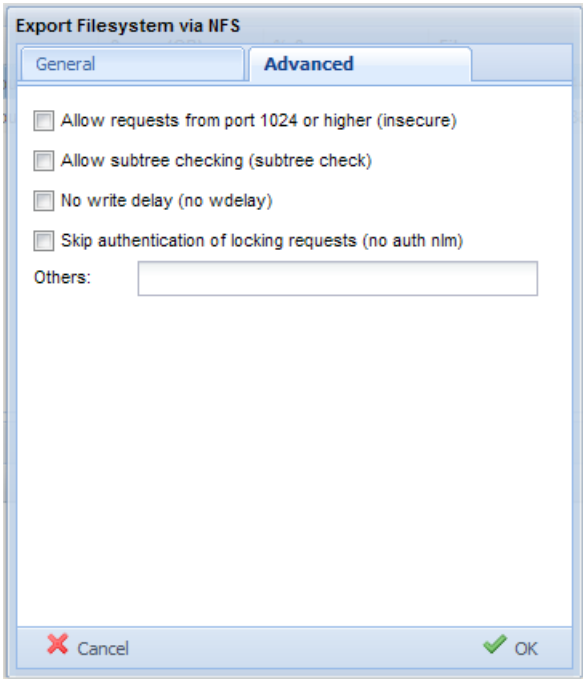

From the CLI

To export a file system from the CLI, use the `ibrix_exportfs` command:

```
<installdirectory>/bin/ibrix_exportfs -f FSNAME -h HOSTNAME -p CLIENT1:PATHNAME1,
CLIENT2:PATHNAME2,.. [-o "OPTIONS"] [-b]
```

The options are as follows:

Option	Description
`-f FSNAME`	The file system to be exported.
`-h HOSTNAME`	The file serving node containing the file system to be exported.
`-p CLIENT1:PATHNAME1, CLIENT2:PATHNAME2,..`	The clients that will access the file system can be a single file serving node, file serving nodes represented by a wildcard, or the world (`:/PATHNAME`). Note that world access omits the client specification but not the colon (for example, `:/usr/src`).
`-o "OPTIONS"`	The default Linux `exportfs` mount options are used unless specific options are provided. The standard NFS export options are supported. Options must be enclosed in double quotation marks (for example, `-o "ro"`). Do not enter an `FSID=` or `sync` option; they are provided automatically.
`-b`	By default, the file system is exported to the NFS client's standby. This option excludes the standby for the file serving node from the export.

For example, to provide NFS clients *.hp.com with read-only access to file system `ifs1` at the directory `/usr/src` on file serving node s1.hp.com:

```
<installdirectory>/bin/ibrix_exportfs -f ifs1 -h s1.hp.com -p *.hp.com:/usr/src -o "ro"
```

To provide world read-only access to file system `ifs1` located at `/usr/src` on file serving node s1.hp.com:

```
<installdirectory>/bin/ibrix_exportfs -f ifs1 -h s1.hp.com -p :/usr/src -o "ro"
```

Unexporting a file system

A file system should be unexported before it is unmounted. On the management console, select the file system, select **NFS Exports** from the lower Navigator pane, and then select **Unexport**.

On the CLI, use the following command:

```
<installdirectory>/bin/ibrix_exportfs -f FSNAME -U -h HOSTNAME -p CLIENT:PATHNAME [-b]
```

Autoconnecting NFS clients

The Autoconnect feature enables NFS clients to mount file systems automatically whenever they are accessed. At the same time, Autoconnect manages how these connections are distributed among file serving nodes. Autoconnect uses the Linux `automount` daemon; working familiarity with `automount` is recommended.

Autoconnect accesses a user-edited script that directs NFS client file requests to the management console, where they are checked against the database and matched to a mount string. The mount string and any mount options are returned to the client, along with the file serving node that the client should use for the mount.

Mount points are stored in the Autoconnect table in the configuration database. Each mount point is described by a user-defined identifier or key, the file system to mount, and any assigned mount options. The `ibrix_autoconnect -l` command displays any current autoconnect entries.

Adding mount points to the autoconnect table

An exported file system must be available when you add mount points. Use the following command to add an autoconnect entry to the table, where *KEY* is a user-defined key, *FSNAME* is the file system name, and *OPTIONS* identifies NFS mount options. (See the Linux `mount` man page for the options.)

```
<installdirectory>/bin/ibrix_autoconnect -A -k KEY -f FSNAME [-o OPTIONS]
```

In the following example, the user-defined key name is `ifs1_rw`. File system `ifs1` will be mounted with options that identify the file system type as NFS and provide read and write permissions for user operations.

```
<installdirectory>/bin/ibrix_autoconnect -A -k ifs1_rw -f ifs1 -o -fstype=nfs,rw
```

Deleting mount points from the autoconnect table

Mount points are deleted by specifying a key entry. Use the following command:

```
<installdirectory>/bin/ibrix_autoconnect -D -k KEY
```

For example, to delete the key `ifs1_ro`:

```
<installdirectory>/bin/ibrix_autoconnect -D -k ifs1_ro
```

Setting up NFS clients

After mount points are defined in the management console, client setup includes editing an Autoconnect script to point to the management console and the correct port, editing the `/etc/auto.master` file on the client, and restarting services.

Before setting up NFS clients, check the following on each client:

- Verify that `automount` is installed by running `/etc/init.d/autofs status`.
- Ensure that `wget` or `curl` is installed if you plan to install a modified `auto.curl` or `auto.wget` script. Use a utility such as `find` to search, if necessary.

Every client must have two files placed in its `/etc` directory. The necessary files are on the management console; edit the files there and copy them to each client. Complete the following steps on each client:

1. Copy the appropriate script (`auto.curl`, `auto.sh`, `auto.wget`) located in *<installdirectory>*`/examples/autoconnect`. Edit the copy to set `fusionmanager` to the IP address of your management console, and verify that the port is set to 9009.

   If you choose, you can write a custom script that provides the same functionality as one of the scripts supplied with X9000 Software. You can also create multiple scripts, each mapping to a different set of primary keys.

2. Set permissions on the script file to make it executable. For example, for a `curl` script, enter the following:

   ```
   chmod +x /etc/auto.curl
   ```

3. Edit the `/etc/auto.master` file to map a base automount mountpoint to the script edited in step 1. For example, to map the base mountpoint `/cluster` to a `curl` script, enter the following:

   ```
   /cluster /etc/auto.curl --timeout=60
   ```

   `timeout` indicates the number of seconds that a connection remains idle before it is automatically disconnected. The mountpoint name could be the same as a key name, but this is not recommended.

4. Copy the edited files to the `/etc` directory on the clients.
5. On each client, restart `autofs` and enable it for future use:

   ```
   /etc/init.d/autofs restart
   chkconfig autofs on
   ```

6. Confirm that `autofs` recognizes the primary keys that you entered in `auto.master`:

   ```
   /etc/init.d/autofs status
   ```

7. Confirm that files are visible at the mountpoint by listing the directory by base mountpoint and key name. For example, for the key name `ifs1_rw`:

   ```
   ls -l /ibrix/ifs1_rw
   ```

6 Configuring authentication for CIFS, FTP, and HTTP

Users accessing CIFS, FTP, or HTTP shares can be authenticated through either Active Directory or Local Users. If you select Active Directory, you can specify the share administrators and enable or disable Linux static user mapping. If you select Local Users, you can create the appropriate local user and local group accounts.

Selecting an authentication method

When selecting your authentication method, you should be aware of the following:

- The Configuration Authentication dialog is used only to select an authentication method. If authentication has already been configured, the dialog does not display the authentication type currently in effect. Instead, that information is displayed on the File Sharing Authentication Settings panel.
- It is possible to configure some servers to use Local Users and other servers to use Active Directory.

To configure authentication on the GUI, select **Cluster Configuration** > **File Sharing Authentication** from the Navigator. The File Sharing Authentication Settings panel shows the current configuration on each server.

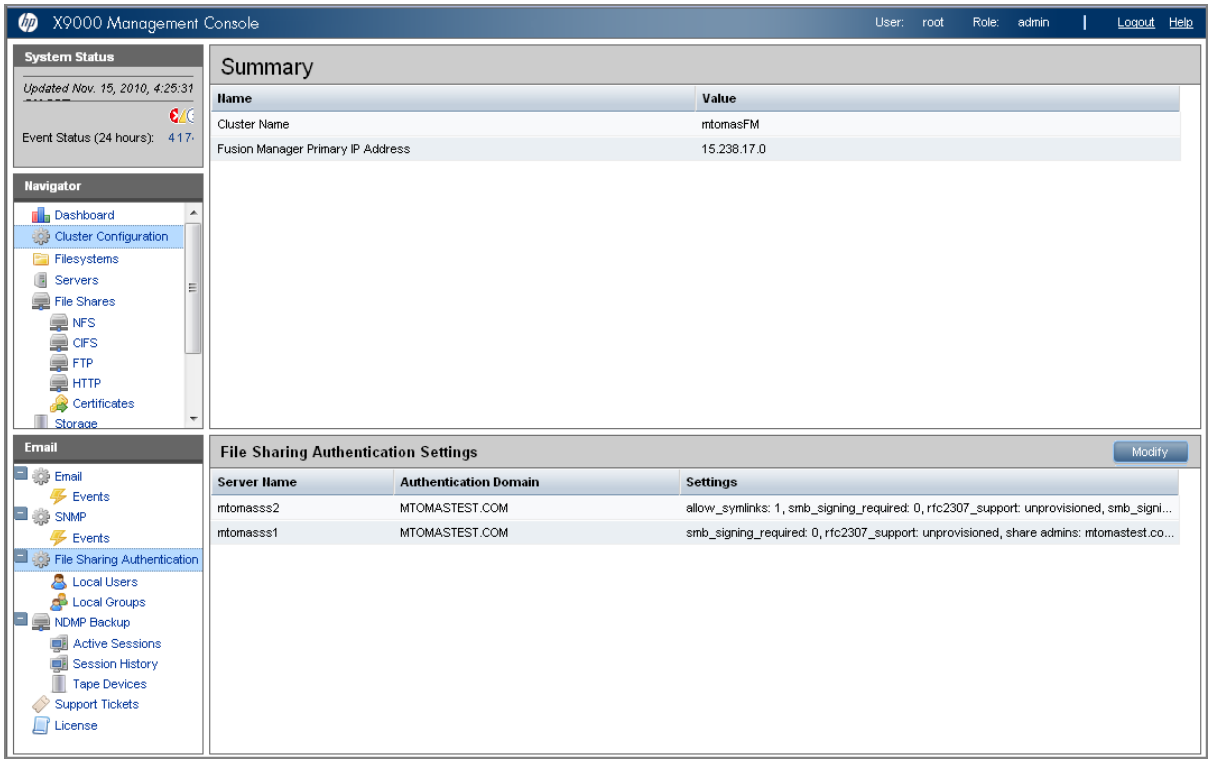

To configure an authentication method, click **Modify** on the File Sharing Authentication Settings panel and complete the Configure File Sharing Authentication dialog. Be sure to click **OK** when you complete the dialog. An authentication method must be in effect when you create shares.

- To configure Active Directory authentication, select a server to configure and enter a list of share administrators (such as `domain\user1, domain\user2, domain\group1`). Select **Active Directory** and enter your domain name, the Auth Proxy username (an AD domain user with privileges to join the specified domain; typically a Domain Administrator), and the password for that user. You can also enable or disable Linux static user mapping. You can apply the configuration just to the selected server, or to all servers.

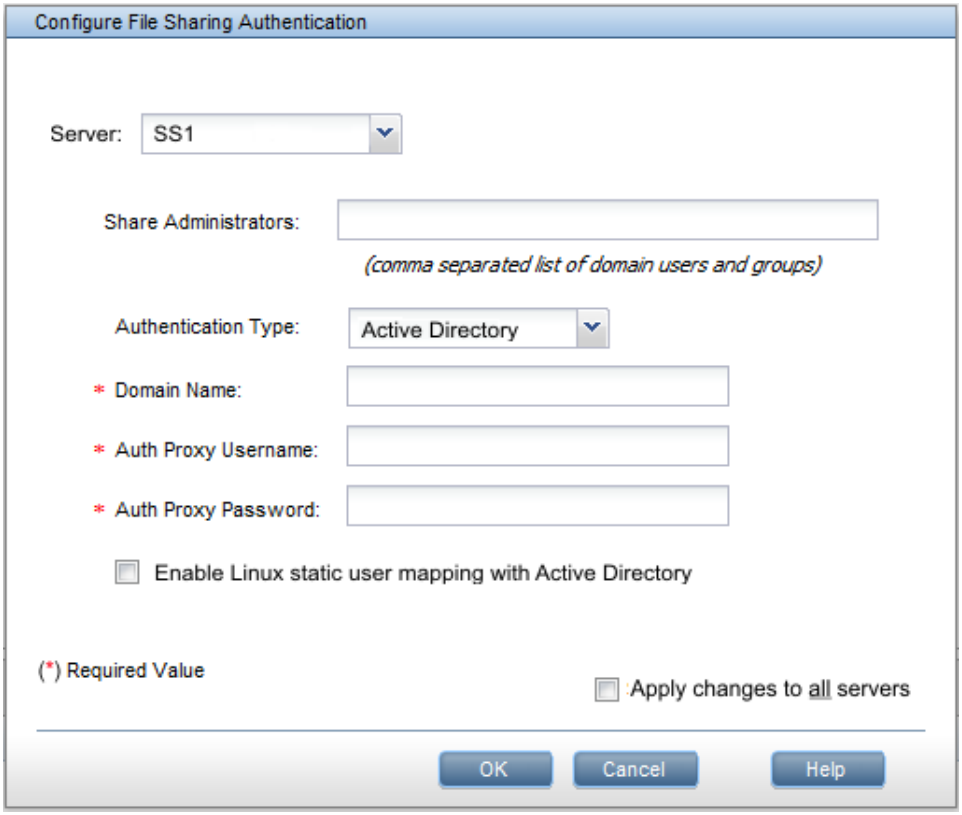

NOTE: When you successfully configure Active Directory authentication, the machine is part of the domain until you remove it from the domain, either with the `ibrix_auth -n` command or with Windows tools. Because Active Directory authentication is a one-time event, it is not necessary to update authentication if you change the proxy user information.

- To configure Local Users authentication, select the server to configure, enter the share administrators (such as `local\root, local\test1, local\test2`), and then select **Local Users**. You can apply the configuration just to the selected server, or to all servers. Local users and groups specified as share administrators must already exist.

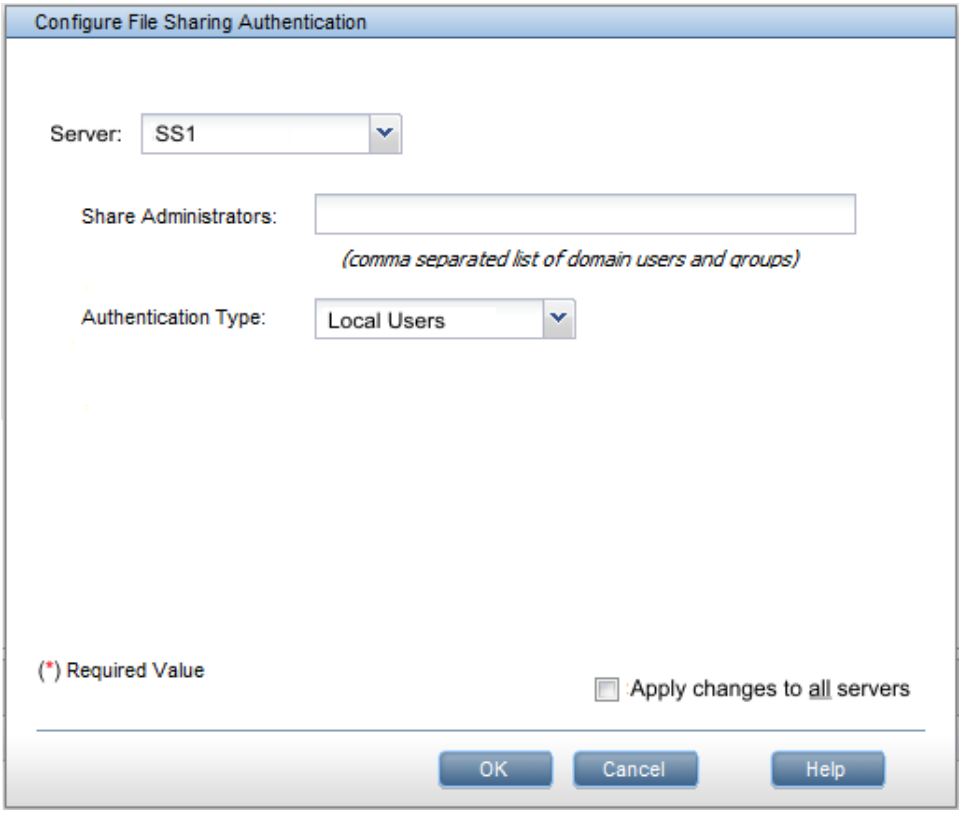

To configure Local Users authentication from the CLI, use the following command:

```
ibrix_auth -N  [-h HOSTLIST]
```

To configure Active Directory authentication, use the following command:

```
ibrix_auth -n DOMAIN_NAME -A AUTH_PROXY_USER_NAME [-P AUTH_PROXY_PASSWORD]
[-S SETTINGLIST] [-h HOSTLIST]
```

RFC2307 is the protocol that enables Linux static user mapping with Active Directory. To enable RFC2307 support from the CLI, use the following command:

```
ibrix_cifsconfig -t [-S SETTINGLIST] [-h HOSTLIST]
```

Enable RFC2307 in the *SETTINGLIST* as follows:

```
rfc2307_support=rfc2307
```

For example:

```
ibrix_cifsconfig -t -S "rfc2307_support=rfc2307"
```

To disable RFC2307, set `rfc2307_support` to `unprovisioned`. For example:

```
ibrix_cifsconfig -t -S "rfc2307_support=unprovisioned"
```

IMPORTANT: After making configuration changes with the `ibrix_cifsconfig -t -S` command, use the following command to restart the CIFS services on all nodes affected by the change.

```
ibrix_server –s –t cifs –c restart [–h SERVERLIST]
```

Clients will experience a temporary interruption in service during the restart.

X9000 management console requirement for Local Users authentication

To use Local Users authentication, your cluster must use the agile management console configuration. Also, the agile management console must be installed on each node participating in CIFS, FTP, or HTTP operations.

IMPORTANT: The servers hosting the agile management console must be file serving nodes. Do not use a dedicated Management Server that cannot be converted to a file serving node as a host for the agile management console.

Complete the following steps to install the agile management console on all file serving nodes:

1. If your cluster is not currently using the agile management console configuration, migrate to the agile configuration using the procedure in the administration guide for your system.
2. Install a passive agile management console on each cluster node that is not currently hosting an active or passive agile management console. You will need the X9000 installation code to perform the installation. As delivered, this code is provided in `/tmp/<system_type>/ibrix`, where `<system_type>` is X9300, X9320, or X9720. If this directory no longer exists, download the installation code from the HP support website for your storage system.

   Run the following command:

   ```
   [root]# <install_code_directory>/ibrixinit -tm -C
   <local_cluster_interface_device> –v <cluster_VIF_IP> -m
   <cluster_netmask> -d <cluster_VIF_device> [-V <user_VIF_IP>] [-N
   <user_netmask>] [-D <user_VIF_device>] -w 9009 –M passive –F
   ```

   For example:

   ```
   [root]# <install_code_directory>/ibrixinit -tm -C bond0 -v 10.30.83.1
   -m 255.255.248.0 -d bond0:0 -V 172.16.3.1 -N 255.255.0.0 -D bond1:0
   -w 9009 -M passive –F
   ```

Configuring local groups and local users

If you selected Local Users as the authentication method, you will need to configure local user and local group accounts. Create a local user account for each user that will be accessing CIFS, FTP, or HTTP shares, and create at least one local group account for the users. The account information is stored internally in the cluster.

When naming local groups and users, you should be aware of the following:

- User and group names must be unique. The new name cannot already be used by another user or group.
- The following names cannot be used for users or groups: `administrator`, `guest`, `root`.

Configuring local group accounts

Local group accounts can be managed from the management console GUI or CLI. On the GUI, select **Cluster Configuration** > **Local Groups** from the Navigator. The Local Groups pane shows the local groups that are currently configured.

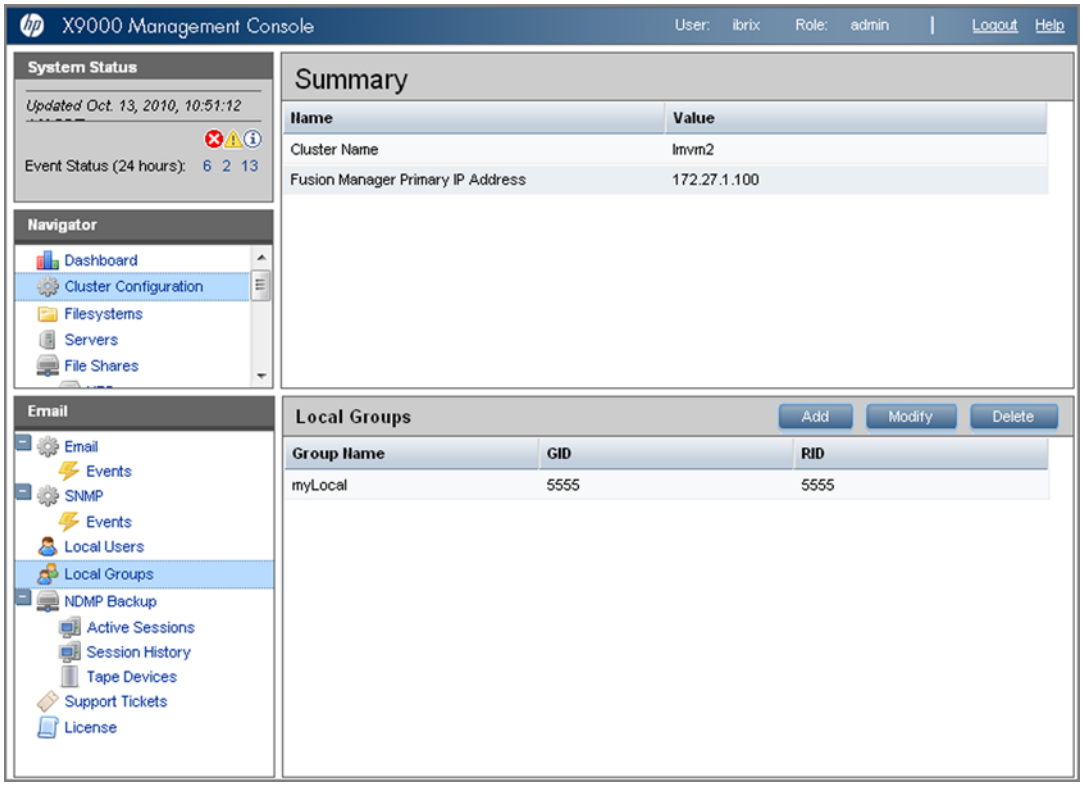

On the CLI, use the following command to view information about all local group accounts:

```
ibrix_localgroups -L
```

To see information for a specific local group account, use the following command:

```
ibrix_localgroups -l -g GROUPNAME
```

Adding a local group

To add a new local group, click **Add** on the Local Groups pane. Then enter the information for the group on the Add Local Group dialog box. The GID and RID will be generated automatically if you do not enter values for them.

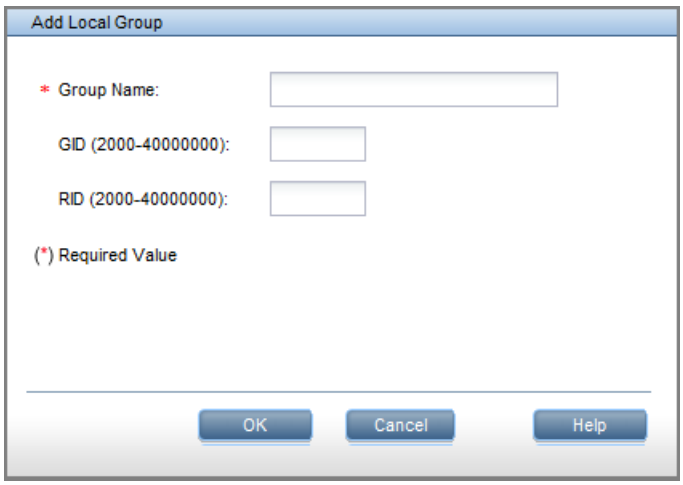

To add a local group account from the CLI, use the following command:

```
ibrix_localgroups -a -g GROUPNAME [-G GROUPID] [-S RID]
```

Modifying a local group

To change the information for a local group account, select the account on the Local Groups pane and click **Modify**. You can then make the necessary changes on the Modify Local Group dialog box. If you are changing the GID or RID for the group, it cannot be the primary group for any local users.

To modify an account from the CLI, use the following command:

```
ibrix_localgroups -m -g GROUPNAME [-G GROUPID] [-S RID]
```

Deleting a local group

To delete a group account, select the account on the Local Groups pane, click **Delete**, and confirm the operation. To delete an account from the CLI, use the following command:

```
ibrix_localgroups -d -g GROUPNAME
```

Configuring local users

Local user accounts can be managed from the management console GUI or CLI. On the GUI, select **Cluster Configuration** > **Local Users**. The Local Users pane shows the local user accounts that are currently configured.

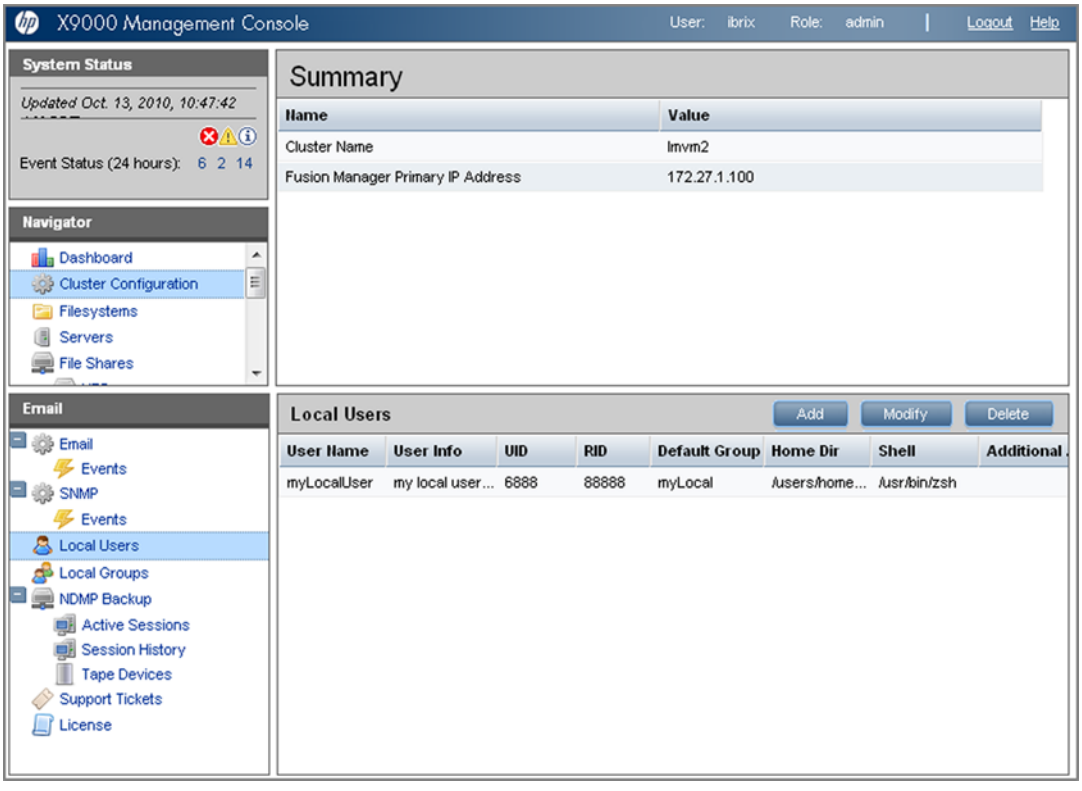

On the CLI, use the following command to view information about all local user accounts:

```
ibrix_localusers -L
```

To see information for a specific local user account, use the following command:

```
ibrix_localusers -l -g USERNAME
```

Adding a local user

To add a new user account, click **Add** on the Local Users pane, and then enter the user's information on the Add Local User dialog. The UID and RID will be generated automatically if you do not enter values for them. The default home directory is `/home/<username>` and the default shell program is `/bin/false`.

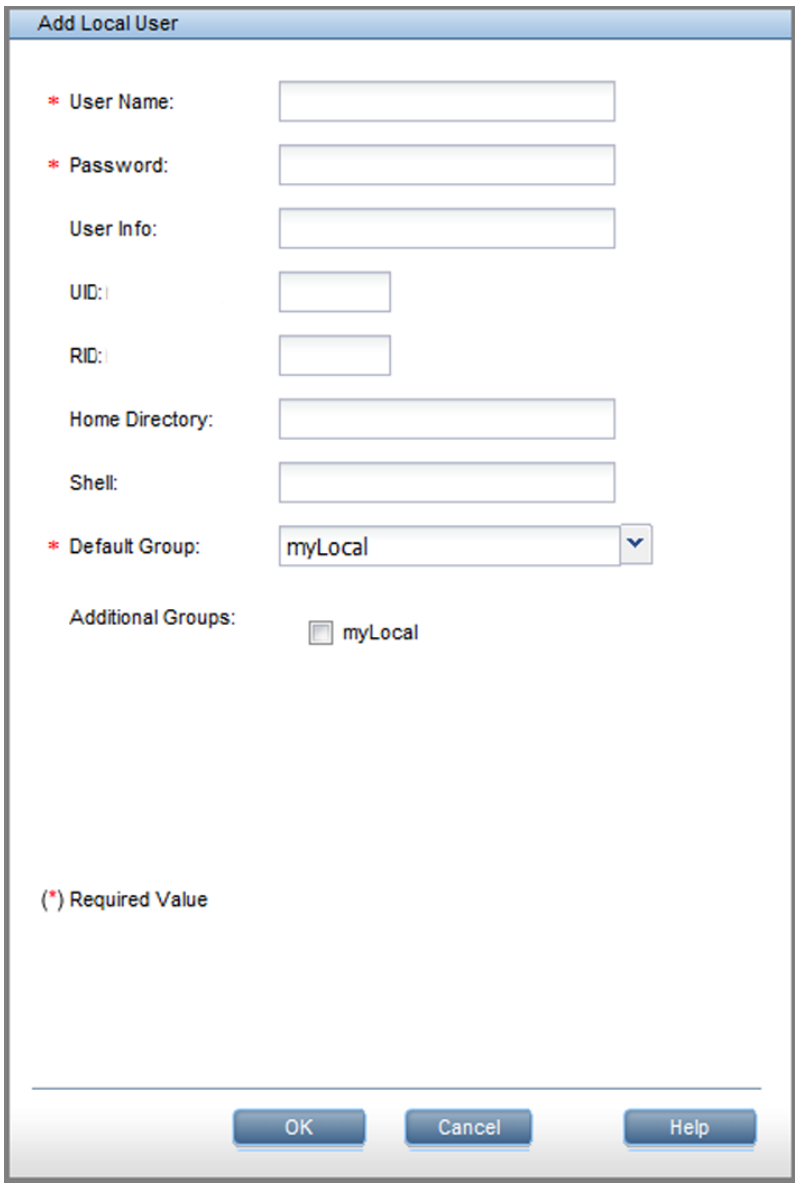

To add a local user account from the CLI, use this command:

```
ibrix_localusers -a -u USERNAME -g DEFAULTGROUP -p PASSWORD [-h HOMEDIR]
[-s SHELL] [-i USERINFO] [-U USERID] [-S RID] [-G GROUPLIST]
```

Modifying a local user

To change the information for a local user account, select the account on the Local Users pane and click **Modify**. You can then make the necessary changes on the Modify Local User dialog box. You cannot change the UID or RID for the account. If it is necessary to change a UID or RID, you will need to delete the account and then recreate it with the new UID or RID.

To modify a local user account from the CLI, use the following command:

```
ibrix_localusers -m -u USERNAME [-g DEFAULTGROUP] [-p PASSWORD] [-h
HOMEDIR] [-s SHELL] [-i USERINFO] [-G GROUPLIST]
```

Deleting a local user

To delete a local user account, select the account on the Local Users pane, click **Delete**, and confirm the operation. To delete an account from the CLI, use the following command:

```
ibrix_localusers -d -u USERNAME
```

7 Using CIFS

The X9000 Software CIFS server implementation allows you to create file shares for data stored on the cluster. The CIFS server provides a true Windows experience for Windows clients. A user accessing a file share on an X9000 system will see the same behavior as on a Windows server.

IMPORTANT: Before configuring CIFS, select an authentication method (either Local Users or Active Directory). See "Configuring authentication for CIFS, FTP, and HTTP" (page 42) for more information.

IMPORTANT: CIFS and X9000 Windows clients cannot be used together because of incompatible AD user to UID mapping. You can use either CIFS or X9000 Windows clients, but not both at the same time.

Configuring file serving nodes for CIFS

To enable file serving nodes to provide CIFS services, you will need to configure the `resolv.conf` file. On each node, the `/etc/resolv.conf` file must include a DNS server that can resolve SRV records for your domain. For example:

```
# cat /etc/resolv.conf

search mycompany.com
nameserver 192.168.100.132
```

To verify that a file serving node can resolve SRV records for your AD domain, run the Linux `dig` command. (In the following example, the Active Directory domain name is `mydomain.com`.)

```
% dig SRV _ldap._tcp.mydomain.com
```

In the output, verify that the `ANSWER SECTION` contains a line with the name of a domain controller in the Active Directory domain. Following is some sample output:

```
; <<>> DiG 9.3.4-P1 <<>> SRV _ldap._tcp.mydomain.com
;; global options:  printcmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 56968
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 2

;; QUESTION SECTION:
;_ldap._tcp.mydomain.com.       IN      SRV

;; ANSWER SECTION:
_ldap._tcp.mydomain.com. 600    IN      SRV     0 100 389 adctrlr.mydomain.com.

;; ADDITIONAL SECTION:
adctrlr.mydomain.com.    3600    IN      A       192.168.11.11

;; Query time: 0 msec
;; SERVER: 192.168.100.132 #53(192.168.100.132)
;; WHEN: Tue Mar 16 09:56:02 2010
;; MSG SIZE  rcvd: 113
```

For more information, see the Linux `resolv.conf(5)` man page.

Starting or stopping the CIFS service and viewing CIFS statistics

IMPORTANT: You will need to start the CIFS service initially on the file serving nodes. Subsequently, the service is started automatically when a node is rebooted.

Use the CIFS tab on the management console GUI to start, stop, or restart the CIFS service on a particular node, or to view CIFS activity statistics for the node. Select **Servers** from the Navigator and then select the appropriate node. Select **CIFS** in the lower Navigator to display the CIFS tab,

which shows CIFS activity statistics on the node. You can start, stop, or restart the CIFS service by clicking the appropriate button.

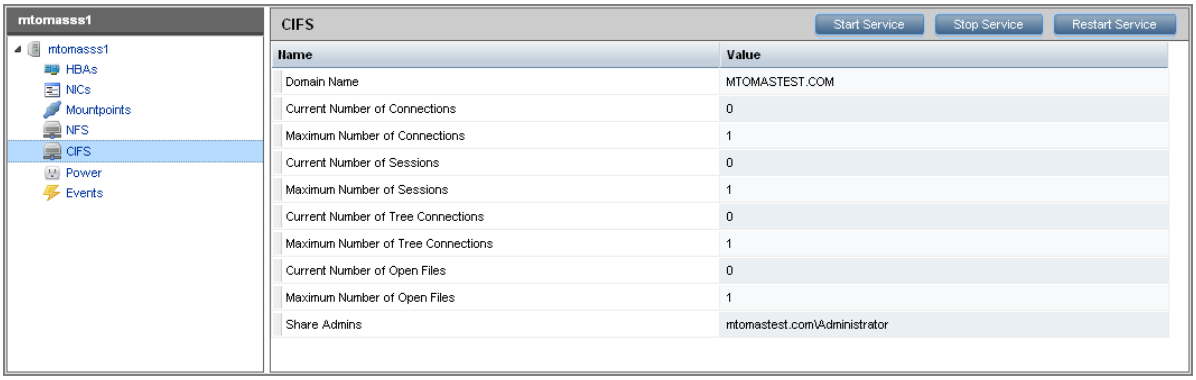

To start, stop, or restart the service from the CLI, use the following command:

```
<installdirectory>/bin/ibrix_server -s -t cifs -c {start|stop|restart}
```

CIFS shares

Windows clients access file systems through CIFS shares. You can use the X9000 management console GUI or CLI to manage shares, or you can use the Microsoft Management Console interface. The CIFS service must be running when you add shares.

X9000 Software supports 5000 shares per server and 3000 connections.

IMPORTANT: The permissions on the directory exporting a CIFS share govern the access rights that are given to the Everyone user as well as to the owner and group of the share. Consequently, the Everyone user may have more access rights than necessary. The administrator should set ACLs on the CIFS share to ensure that users have only the appropriate access rights. Alternatively, permissions can be set more restrictively on the directory exporting the CIFS share.

Managing CIFS shares with the X9000 management console GUI or CLI

To view existing CIFS shares on the GUI, select **File Shares** > **CIFS** from the Navigator. The CIFS Shares window shows the file system being shared, the hosts (or servers) providing access, the name of the share, the export path, and the options applied to the share.

NOTE: When `externally managed` appears in the option list for a share, that share is being managed with the Microsoft Management Console interface. The management conosle GUI or CLI cannot be used to change the permissions for the share.

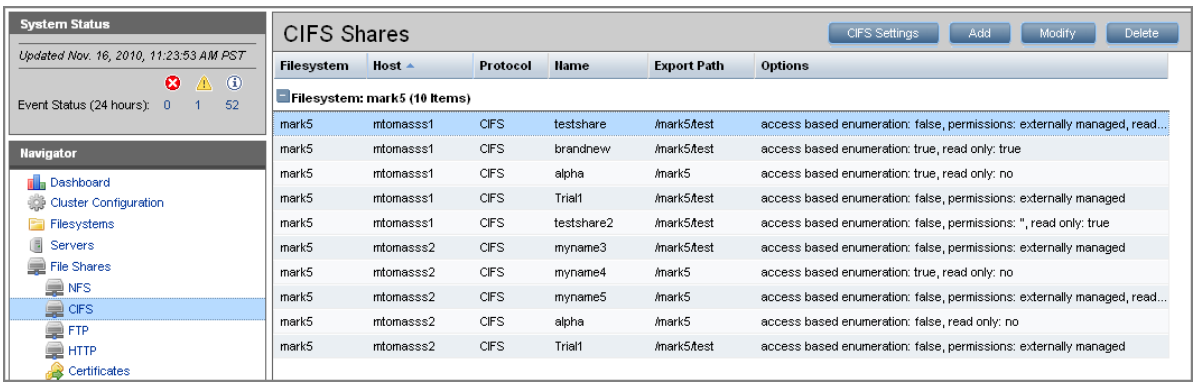
System Status
Updated Nov. 16, 2010, 11:23:53 AM PST
Event Status (24 hours): 0 1 52
Navigator
Dashboard
Cluster Configuration
Filesystems
Servers
File Shares
NFS
CIFS
FTP
HTTP
Certificates
CIFS Shares
CIFS Settings
Add
Modify
Delete
Filesystem
Host
Protocol
Name
Export Path
Options
Filesystem: mark5 (10 Items)
mark5 mtomasss1 CIFS testshare /mark5/test access based enumeration: false, permissions: externally managed, read...
mark5 mtomasss1 CIFS brandnew /mark5/test access based enumeration: true, read only: true
mark5 mtomasss1 CIFS alpha /mark5 access based enumeration: true, read only: no
mark5 mtomasss1 CIFS Trial1 /mark5/test access based enumeration: false, permissions: externally managed
mark5 mtomasss1 CIFS testshare2 /mark5/test access based enumeration: false, permissions: ", read only: true
mark5 mtomasss2 CIFS myname3 /mark5/test access based enumeration: false, permissions: externally managed
mark5 mtomasss2 CIFS myname4 /mark5 access based enumeration: true, read only: no
mark5 mtomasss2 CIFS myname5 /mark5 access based enumeration: false, permissions: externally managed, read...
mark5 mtomasss2 CIFS alpha /mark5 access based enumeration: false, read only: no
mark5 mtomasss2 CIFS Trial1 /mark5/test access based enumeration: false, permissions: externally managed

From the CIFS Shares window, you can add, modify, or delete shares, and you can configure global CIFS settings. (Currently, the only global CIFS setting is SMB signing.)

You can also view CIFS shares for a specific file system. Select that file system on the GUI, and then select CIFS Shares from the lower Navigator.

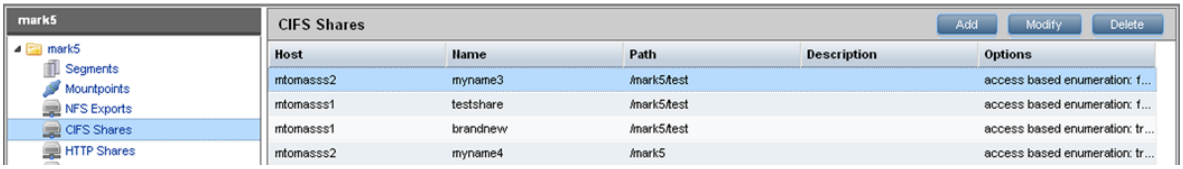

You can add, modify, or delete shares from this window, but you cannot configure CIFS settings, as those settings apply to all CIFS shares configured in the cluster.

To view CIFS shares using the CLI, execute the following command:

```
<installdirectory>/bin/ibrix_cifs -i [-h HOSTLIST]
```

Configuring SMB signing

Use the SMB signing feature to specify whether clients must support SMB signing to access CIFS shares. You can configure SMB signing as follows:

- **Disabled.** SMB signing is not in effect.
- **Enabled.** Both clients that support SMB signing and clients that do not support SMB clients can connect to CIFS shares.
- **Enabled and required.** Only those clients supporting SMB signing can connect to CIFS shares. This is the default setting.

When configuring SMB signing, you should be aware of the following:

- The Configure CIFS Protocol Settings dialog is used only to enable or disable SMB signing. It does not display whether SMB signing is currently enabled or disabled.
- SMB signing must not be required to support connections from 10.5 and 10.6 Mac clients.
- It is possible to configure SMB signing differently on individual servers. Backup CIFS servers should have the same settings to ensure that clients can connect after a failover.
- The SMB signing settings specified here are not affected by Windows domain group policy settings when joined to a Windows domain.

To configure SMB signing, select **File Shares** > **CIFS** from the Navigator, and then click **CIFS Settings** on the CIFS Shares window. Select a server on the Configure CIFS Protocol Settings dialog box and apply the appropriate setting just to that server or to all servers.

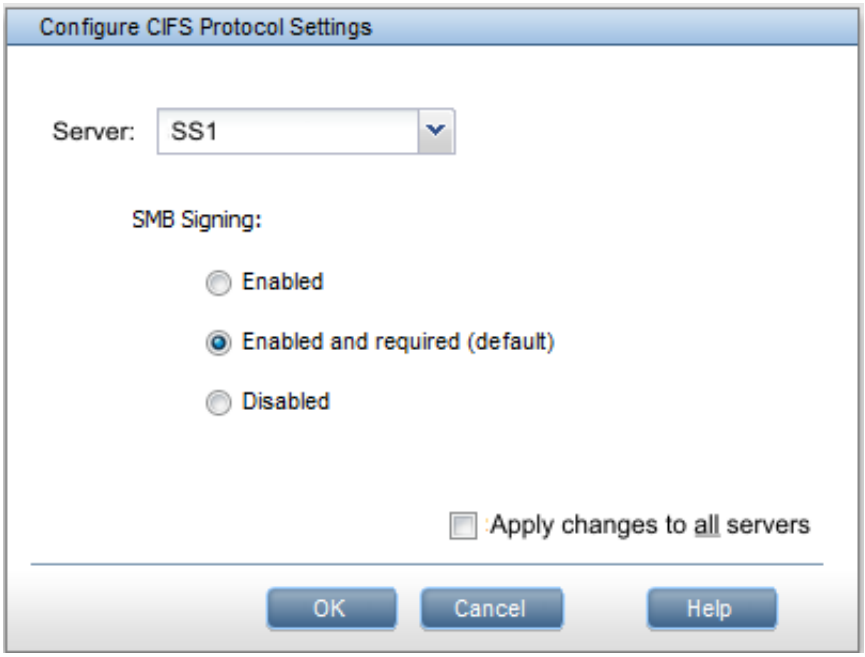

To view the current setting for SMB signing, use the following command:

```
ibrix_cifsconfig -i
```

To configure SMB signing from the command line, use the following command:

```
ibrix_cifsconfig -t -S SETTINGLIST
```

You can specify the following values in the *SETTINGLIST*:

```
smb signing enabled
smb signing required
```

Use commas to separate the settings, and enclose the list in quotation marks. For example, the following command sets SMB signing to enabled and required:

```
ibrix_cifsconfig –t –S "smb signing enabled=1,smb signing required=1"
```

To disable SMB signing, enter *settingname=* with no value. For example:

```
ibrix_cifsconfig –t –S "smb signing enabled=,smb signing required="
```

IMPORTANT: After making configuration changes with the `ibrix_cifsconfig -t -S` command, use the following command to restart the CIFS services on all nodes affected by the change.

```
ibrix_server –s –t cifs –c restart [–h SERVERLIST]
```

Clients will experience a temporary interruption in service during the restart.

Adding a CIFS share

To add CIFS shares from the GUI, click **Add** on either CIFS Shares window. (You may need to select the file system for the share on the Add File Share dialog box.)

The Add a CIFS Share dialog box allows you to share the entire filesystem or a specific subdirectory. Enter a name and description for the share, select the appropriate permissions, and select the servers on which the share will be created. Note the following:

- Do not include any of the following special characters in a share name. If the name contains any of these special characters, the share might not be set up properly on all nodes in the cluster.

  ```
  ' & ( [ { $ ` , / \
  ```

- Do not include any of the following special characters in the share description. If a description contains any of these special characters, the description might not propagate correctly to all nodes in the cluster.

  ```
  * % + & `
  ```

The Access Based Enumeration option allows users to see only the files and folders to which they have access on the file share. Select this option if appropriate for the share.

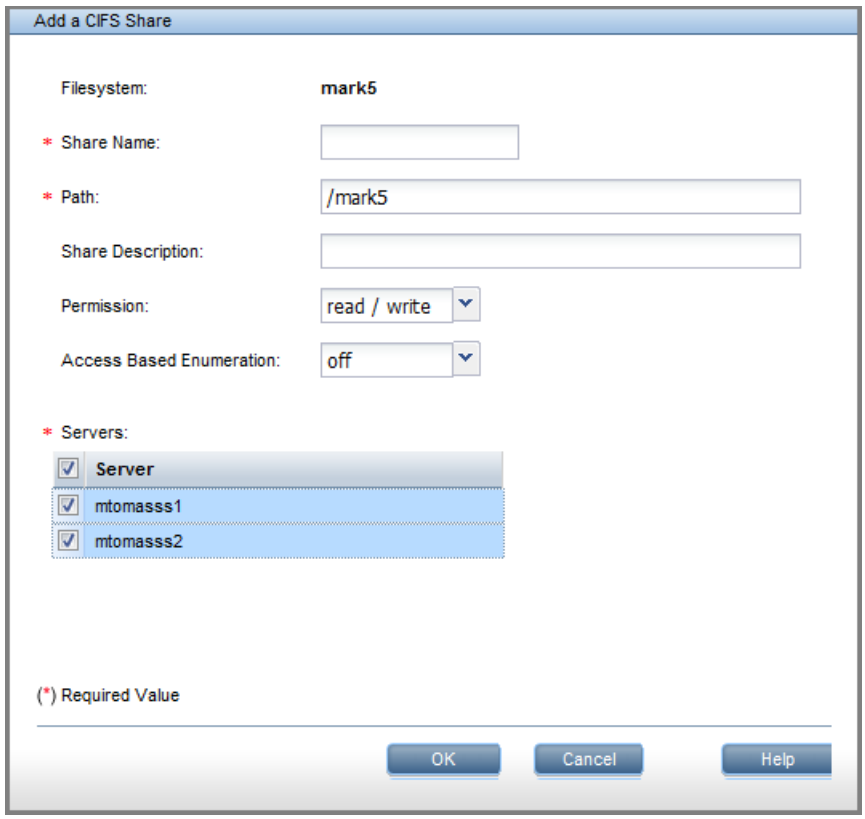

To add shares from the CLI, use the following command:

```
<installdirectory>/bin/ibrix_cifs -a -f FSNAME –s SHARENAME –p SHAREPATH
[-S SETTINGLIST] [-h HOSTLIST]
```

To see the valid settings for the `-S` option, use the following command:

```
<installdirectory>/bin/ibrix_cifs -L
```

NOTE: Be sure to use the `ibrix_cifs` command located in `<installdirectory>/bin`. The `ibrix_cifs` command located in `/usr/local/bin/init` is used internally by X9000 Software and should not be run directly.

Modifying a CIFS share

To change the configuration of a CIFS share, select the share from either CIFS Shares window and click **Modify**. You can change the share description, the permissions, and whether Access Based Enumeration is enabled.

To modify a share using the CLI, execute the following command:

```
<installdirectory>/bin/ibrix_cifs -m -s SHARENAME [-S SETTINGLIST] [-h HOSTLIST]
```

Deleting a CIFS share

To delete a CIFS share, select the share from either CIFS Shares window, click **Delete**, and confirm the operation.

To delete a CIFS share using the CLI, execute the following command:

```
<installdirectory>/bin/ibrix_cifs –d -s SHARENAME [-h HOSTLIST]
```

Managing CIFS shares with Microsoft Management Console

The Microsoft Management Console (MMC) can be used to add, view, or delete CIFS shares. Administrators running MMC must have X9000 Software share management privileges.

NOTE: To use MMC to manage CIFS shares, you must be authenticated as a user with share modification permissions.

Connecting to cluster nodes

When connecting to cluster nodes, use the procedure corresponding to the Windows operating system on your machine.

Windows XP, Windows 2003 R2:

Complete the following steps:

1. Open the Start menu, select **Run**, and specify **mmc** as the program to open.
2. On the Console Root window, select **File** > **Add/Remove Snap-in**.
3. On the Add/Remove Snap-in window, click **Add.**
4. On the Add Standalone Snap-in window, select **Shared Folders** and click **Add**.
5. On the Shared Folders window, select **Another computer** as the computer to be managed, enter or browse to the computer name, and click **Finish**.

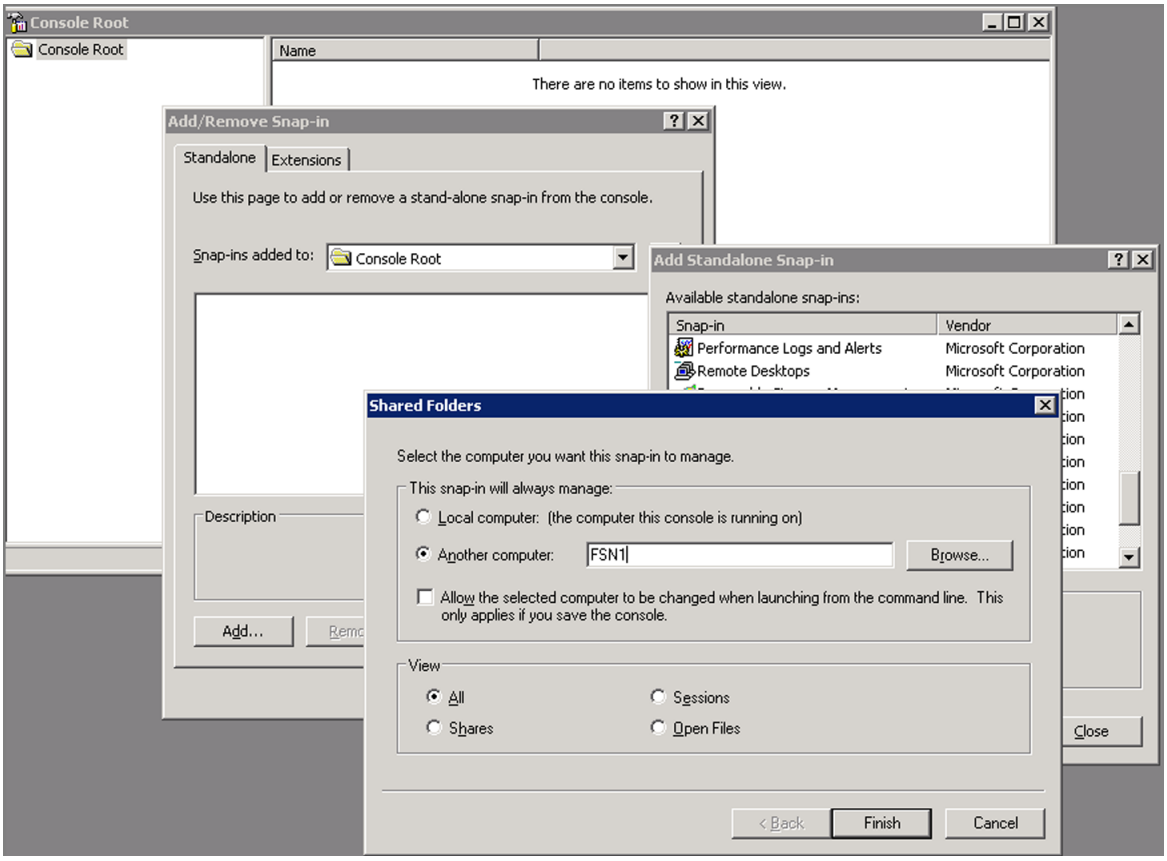

6. Click **Close** > **OK** to exit the dialogs.
7. Expand **Shared Folders (*<address>*)**.
8. Select **Shares** and manage the shares as needed.

Windows Vista, Windows 2008, Windows 7:

Complete the following steps:

1. Open the Start menu and enter **mmc** in the **Start Search** box. You can also enter **mmc** in a DOS cmd window.
2. On the User Account Control window, click **Continue**.
3. On the Console 1 window, select **File** > **Add/Remove Snap-in**.
4. On the Add or Remove Snap-ins window, select **Shared Folders** and click **Add**.
5. On the Shared Folders window, select **Another computer** as the computer to be managed, enter or browse to the computer name, and click **Finish**.

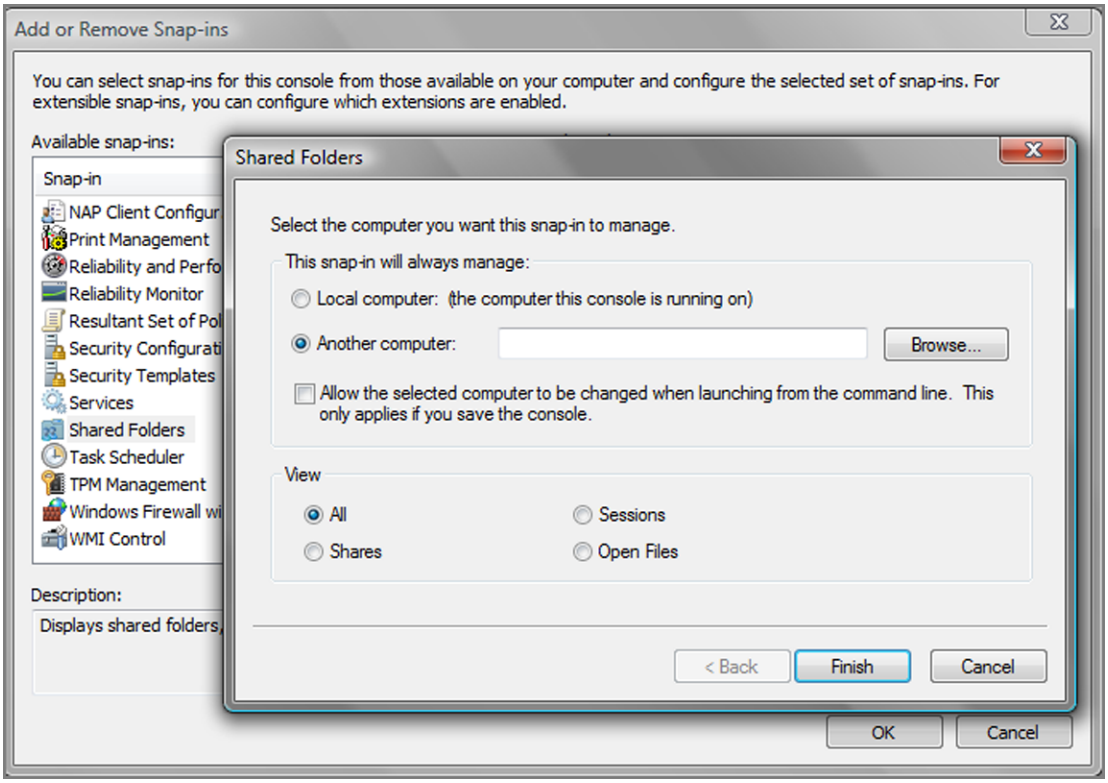

6. Click **OK** to exit the Add or Remove Snap-ins window.
7. Expand **Shared Folders (*<address>*)**.
8. Select **Shares** and manage the shares as needed.

Saving MMC settings

You can save your MMC settings to use when managing shares on this server in later sessions. Complete these steps:

1. On the MMC, select **File** > **Save As**.
2. Enter a name for the file. The name must have the suffix `.msc`.
3. Select **Desktop** as the location to save the file, and click **Save**.
4. Select **File** > **Exit**.

Granting share management privileges

Use the following command to grant administrators X9000 Software share management privileges. The users you specify must already exist. Be sure to enclose the user names in square brackets.

```
ibrix_auth -t -S 'share admins=[domainname\username,domainname\username]'
```

The following example gives share management privileges to a single user:

```
ibrix_auth -t -S 'share admins=[domain\user1]'
```

If you specify multiple administrators, use commas to separate the users. For example:

```
ibrix_auth -t -S 'share admins=[domain\user1, domain\user2,
domain\user3]'
```

Adding CIFS shares

CIFS shares can be added with the MMC, using the share management plug-in. When adding shares, you should be aware of the following:

- The share path must include the X9000 file system name. For example, if the file system is named `data`, you could specify `C:\fs1\folder1`.

  NOTE: The Browse button cannot be used to locate the file system.

- The directory to be shared will be created if it does not already exist.
- The permissions on the shared directory will be set to 777. It is not possible to change the permissions on the share.
- Do not include any of the following special characters in a share name. If the name contains any of these special characters, the share might not be set up properly on all nodes in the cluster.

  ```
  ' & ( [ { $ ` , / \
  ```

- Do not include any of the following special characters in the share description. If a description contains any of these special characters, the description might not propagate correctly to all nodes in the cluster.

  ```
  * % + & `
  ```

- The management console GUI or CLI cannot be used to alter the permissions for shares created or managed with Windows Share Management. The permissions for these shares are marked as "externally managed" on the GUI and CLI.

Open the MMC with the Shared Folders snap-in that you created earlier. On the Select Computer dialog box, enter the IP address of a server that will host the share.

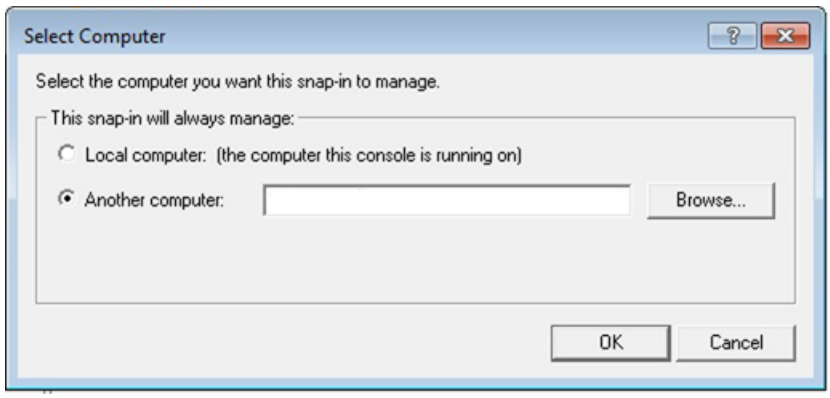

The Computer Management window shows the shares currently available from server.

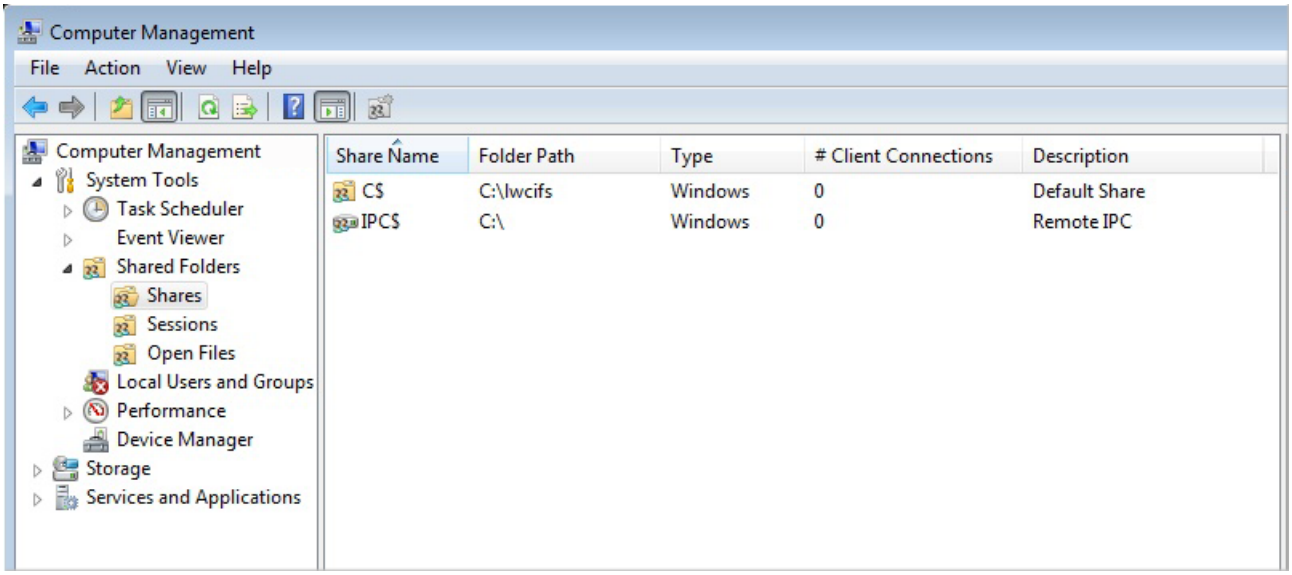

To add a new share, select **Shares** > **New Share** and run the Create A Shared Folder Wizard. On the Folder Path page, enter the path to the share, being sure to include the file system name.

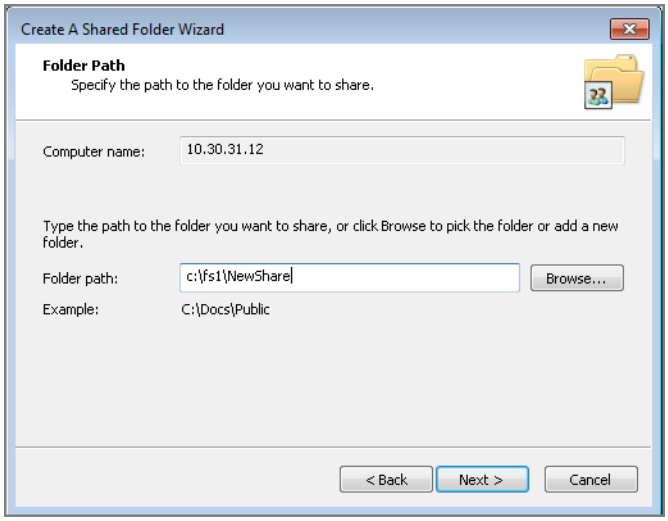
Create A Shared Folder Wizard
Folder Path
Specify the path to the folder you want to share.
Computer name:
10.30.31.12
Type the path to the folder you want to share, or click Browse to pick the folder or add a new folder.
Folder path:
c:\fs1\NewShare
Browse...
Example:
C:\Docs\Public
< Back
Next >
Cancel

When you complete the wizard, the new share appears on the Computer Management window.

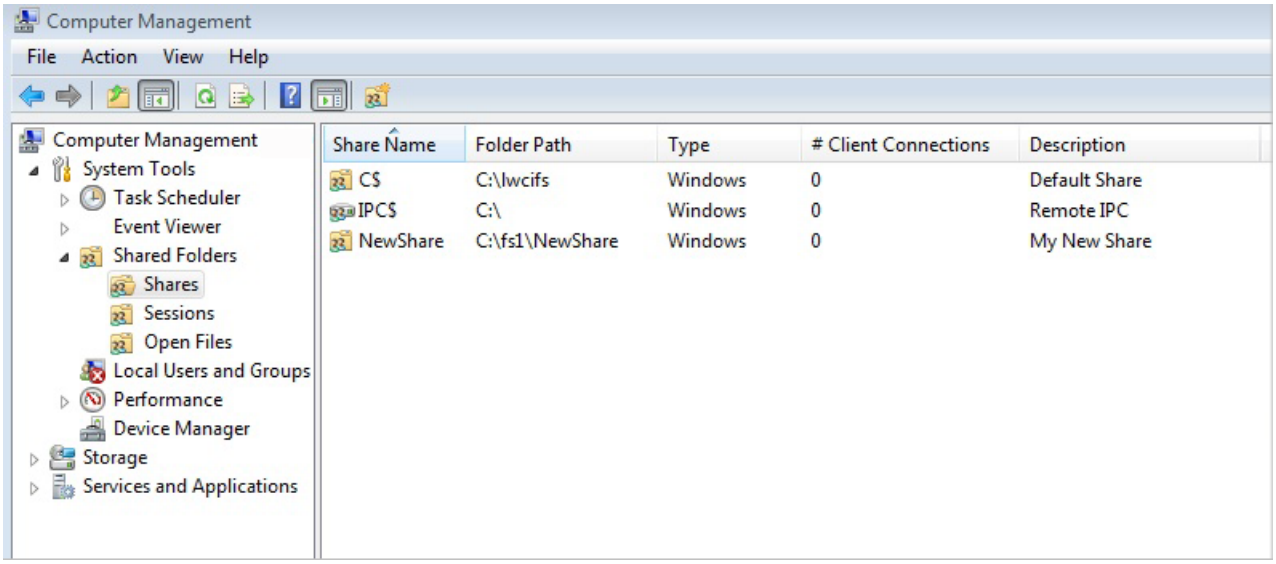

Deleting CIFS shares

To delete a CIFS share, select the share on the Computer Management window, right-click, and select **Delete**.

Linux static user mapping with Active Directory

Linux static user mapping (also called UID/GID mapping or RFC2307 support) allows you to use LDAP as a Network Information Service. When this feature is enabled, you can assign UIDs, GIDs, and other POSIX attributes such as the home directory, primary group and shell to users and groups in Active Directory.

NOTE: Before configuring Linux static user mapping, enable the feature through the management console or CLI. See "Configuring authentication for CIFS, FTP, and HTTP" (page 42) for more information.

To use Linux static user mapping, complete these steps:

- Configure Active Directory.
- Assign POSIX attributes to users and groups in Active Directory.

NOTE: Mapping UID 0 and GID 0 to any AD user or group is not compatible with CIFS static mapping.

Configuring Active Directory

Your Windows Domain Controller machines must be running Windows Server 2003 R2 or Windows Server 2008 R2. Configure the Active Directory domain as follows:

- Install Identity Management for UNIX.
- Activate the Active Directory Schema MMC snap-in.
- Add the `uidNumber` and `gidNumber` attributes to the partial-attribute-set of the AD global catalog.

You can perform these procedures from any domain controller. However, the account used to add attributes to the partial-attribute-set must be a member of the `Schema Admins` group.

Installing Identity Management for UNIX

To install Identity Management for UNIX on a domain controller running Windows Server 2003 R2, see the following Microsoft TechNet Article:

http://technet.microsoft.com/en-us/library/cc778455(WS.10).aspx

To install Identity Management for UNIX on a domain controller running Windows Server 2008 R2, see the following Microsoft TechNet article:

http://technet.microsoft.com/en-us/library/cc731178.aspx

Activating the Active Directory Schema MMC snap-in

Use the Active Directory Schema MMC snap-in to add the attributes. To activate the snap-in, complete the following steps:

1. Click **Start**, click **Run**, type **mmc**, and then click **OK**.
2. On the MMC Console menu, click **Add/Remove Snap-in**.
3. Click **Add**, and then click **Active Directory Schema**.
4. Click **Add**, click **Close**, and then click **OK**.

Adding uidNumber and gidNumber attributes to the partial-attribute-set

To make modifications using the Active Directory Schema MMC snap-in, complete these steps:

1. Click the **Attributes** folder in the snap-in.
2. In the right pane, scroll to the desired attribute, right-click the attribute, and then click **Properties**.

   Select **Replicate this attribute to the Global Catalog**, and click **OK**.

The following screen shows the properties for the `uidNumber` attribute:

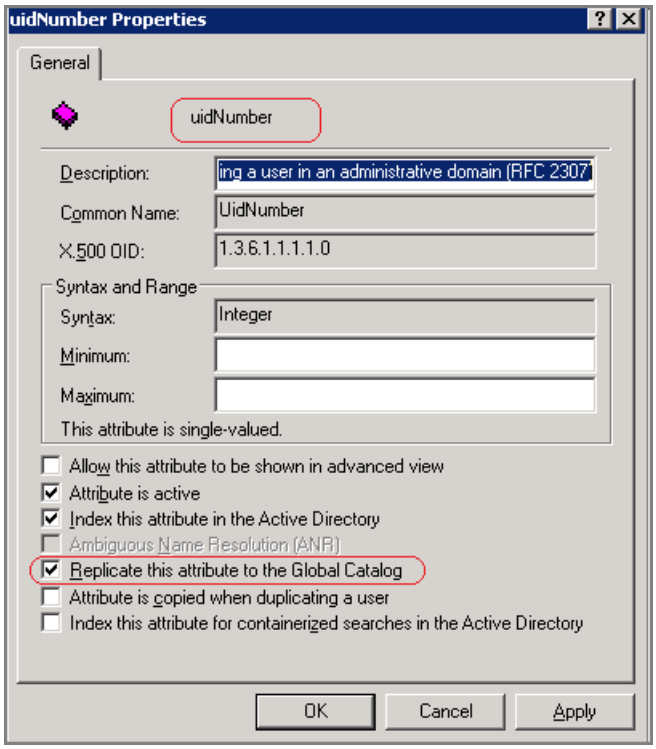

The next screen shows the properties for the `gidNumber` attribute.

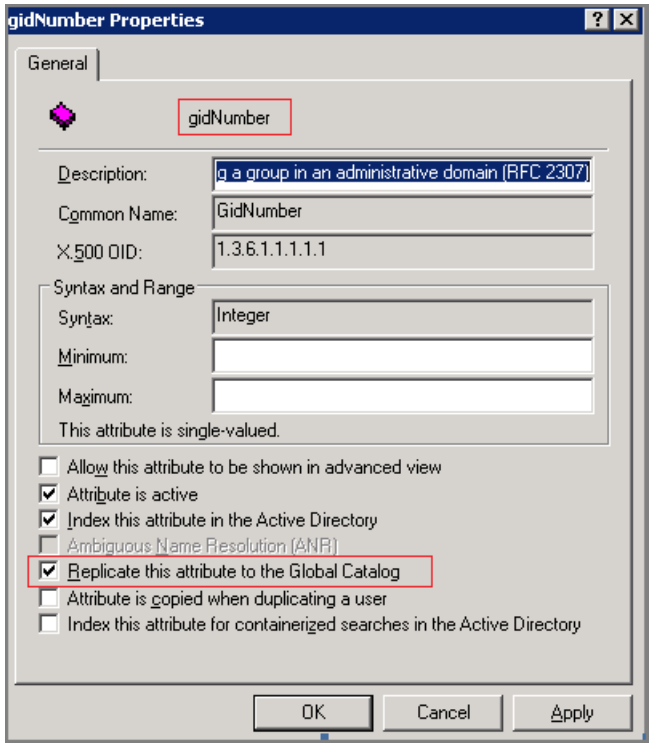

The following article provides more information about modifying attributes in the Active Directory global catalog:

http://support.microsoft.com/kb/248717

Assigning attributes

To set POSIX attributes for users and groups, start the **Active Directory Users and Computers** GUI on the Domain Controller. Open the **Administrator Properties** dialog box, and go to the **UNIX Attributes** tab. For users, you can set the UID, login shell, home directory, and primary group. For groups, set the GID.

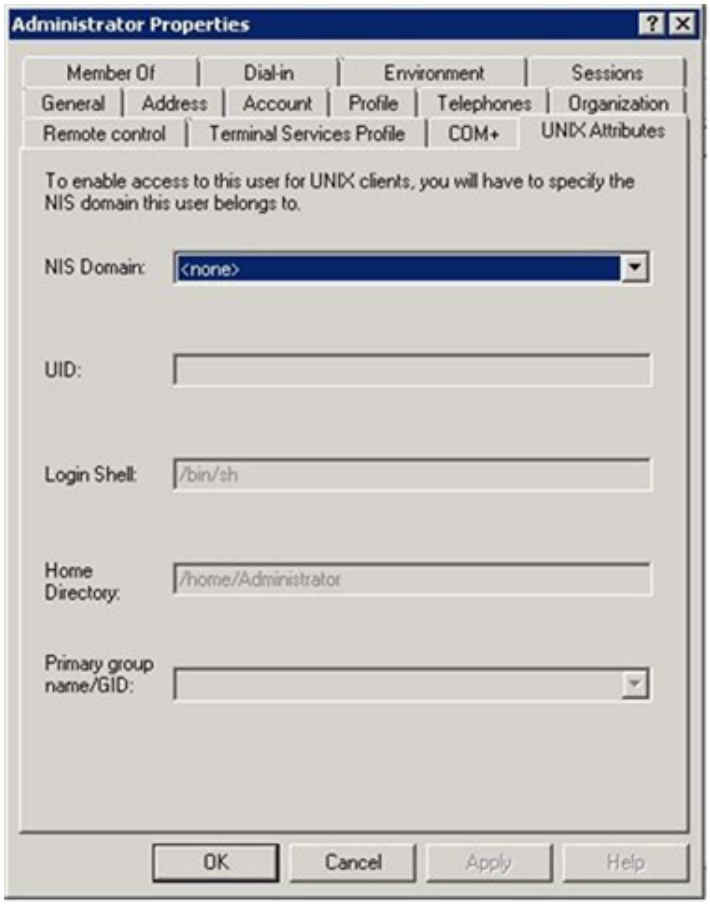

Consolidating SMB servers with common share names

If your SMB servers previously used the same share names, you can consolidate the servers without changing the share name requested on the client side. For example, you might have three SMB servers, `SRV1`, `SRV2`, and `SRV3`, that each have a share named `DATA`. `SRV3` points to a shared drive that has the same path as `\\SRV1\DATA`; however, users accessing `SRV3` have different permissions on the share.

To consolidate the three servers, we will take these steps:

1. Assign Vhost names `SRV1`, `SRV2`, and `SRV3`.
2. Create virtual interfaces (VIF) for the IP addresses used by the servers. For example, Vhost `SRV1` has VIF 99.10.10.101 and Vhost `SRV2` has VIF 99.10.10.102.
3. Map the old share names to new share names. For example, map `\\SRV1\DATA` to new share `srv1-DATA`, map `\\SRV2\DATA` to new share `srv2-DATA`, and map `\\SRV3\DATA` to `srv3-DATA`.
4. Create the new shares on the cluster storage and assign each share the appropriate path. For example, assign `srv1-DATA` to `/srv1/data`, and assign `srv2-DATA` to `/srv2/data`.

   Because `SRV3` originally pointed to the same share as `SRV1`, we will assign the share `srv3-DATA` the same path as `srv1-DATA`, but set the permissions differently.
5. Optionally, create a share having the original share name, `DATA` in our example. Assign a path such as `/ERROR/DATA` and place a file in it named `SHARE_MAP_FAILED`. Doing this ensures that if a user configuration error occurs or the map fails, clients will not gain access to the wrong shares. The file name notifies the user that their access has failed.

When this configuration is in place, a client request to access share `\\srv1/data` will be translated to share `srv1-DATA` at `/srv1/data` on the filesystem. Client requests for `\\srv3/data` will

also be translated to `/srv1/data`, but the clients will have different permissions. The client requests for `\\srv2/data` will be translated to share `srv2-DATA` at `/srv2/data`.

Client utilities such as `net use` will report the requested share name, not the new share name.

Mapping old share names to new share names

Mappings are defined in the `/etc/likewise/vhostmap` file. Use a text editor to create and update the file. Each line in the file contains a mapping in the following format:

```
VIF (or VhostName)|oldShareName|newShareName
```

If you enter a `VhostName`, it will be changed to a VIF internally.

The `oldShareName` is the user-requested share name from the client that needs to be translated into a unique name. This unique name (the `newShareName`) is used when establishing a mount point for the share.

Following are some entries from a `vhostmap` file:

```
99.30.8.23|salesd|q1salesd
99.30.8.24|salesd|q2salesd
salesSrv|salesq|q3salesd
```

When editing the `/etc/likewise/vhostmap` file, note the following:

- All `VIF|oldShareName` pairs must be unique.
- The following characters cannot be used in a share name: `" / \ | [ ] < > + : ; , ? * =`
- Share names are case insensitive, and must be unique with respect to case.
- The `oldShareName` and `newShareName` do not need to exist when creating the file; however, they must exist for a connection to be established to the share.
- If a client specifies a share name that is not in the file, the share name will not be translated.
- Care should be used when assigning share names longer than 12 characters. Some clients impose a limit of 12 characters for a share name.
- Verify that the IP addresses specified in the file are legal and that Vhost names can be resolved to an IP address. IP addresses must be IP4 format, which limits the addresses to 15 characters.

IMPORTANT: When you update the `vhostmap` file, the changes take effect a few minutes after the map is saved. If a client attempts a connection before the changes are in effect, the previous map settings will be used. To avoid any delays, make your changes to the file when the CIFS service is down.

After creating or updating the `vhostmap` file, copy the file manually to the other servers in the cluster.

CIFS clients

CIFS clients access shares on the X9000 Software cluster in the same way they access shares on a Windows server.

Differences in locking behavior

When CIFS clients access a share from different servers, as in the X9000 Software environment, the behavior of byte-range locks differs from the standard Windows behavior, where clients access a share from the same server. You should be aware of the following:

- Zero-length byte-range locks acquired on one file serving node are not observed on other file serving nodes.
- Byte-range locks acquired on one file serving node are not enforced as mandatory on other file serving nodes.
- If a shared byte-range lock is acquired on a file opened with write-only access on one file serving node, that byte-range lock will not be observed on other file serving nodes. ("Write-only access" means the file was opened with GENERIC_WRITE but not GENERIC_READ access.)
- If an exclusive byte-range lock is acquired on a file opened with read-only access on one file serving node, that byte-range lock will not be observed on other file serving nodes. ("Read-only access" means the file was opened with GENERIC_READ but not GENERIC_WRITE access.)

Permissions in a cross-protocol CIFS environment

The manner in which the CIFS server handles permissions affects the use of files by both Windows and Linux clients. Following are some considerations.

How the CIFS server handles UIDs and GIDs

The CIFS server provides a true Windows experience for Windows users. Consequently, it must be closely aligned with Windows in the way it handles permissions and ownership on files. Windows uses ACLs to control permissions on files. The CIFS server puts a bit-for-bit copy of the ACLs on the Linux server (in the files on the X9000 file system), and validates file access through these permissions.

ACLs are tied to Security Identifiers (SIDs) that uniquely identify users in the Windows environment, and which are also stored on the file in the Linux server as a part of the ACLs. SIDs are obtained from the authenticating authority for the Windows client (in X9000 Software, an Active Directory server). However, Linux does not understand Windows-style SIDs; instead, it has its own permissions control scheme based on UID/GID and permissions bits (mode bits, sticky bits). Since this is the native permissions scheme for Linux, the CIFS server must make use of it to access files on behalf of a Windows client; it does this by mapping the SID to a UID/GID and impersonating that UID/GID when accessing files on the Linux file system.

From a Windows standpoint, all of the security for the X9000 Software-resident files is self-consistent; Windows clients understand ACLs and SIDs, and understand how they work together to control access to and security for Windows clients. The CIFS server maintains the ACLs as requested by the Windows clients, and emulates the inheritance of ACLs identically to the way Windows servers maintain inheritance. This creates a true Windows experience around accessing files from a Windows client.

This mechanism works well for pure Linux environments, but (like the CIFS server) Linux applications do not understand any permissions mechanisms other than their own. Note that a Linux application can also use POSIX ACLs to control access to a file; POSIX ACLs are honored by the CIFS server, but will not be inherited or propagated. The CIFS server also does not map POSIX ACLs to be compatible with Windows ACLs on a file.

These permission mechanisms have some ramifications for setting up shares, and for cross-protocol access to files on an X9000 system. The details of these ramifications follow.

Permissions, UIDs/GIDs, and ACLs

The X9000 Software CIFS server does not attempt to maintain two permission/access schemes on the same file. The CIFS server is concerned with maintaining ACLs, so performs ACL inheritance

and honors ACLS. The UID/GIDs and permission bits for files on a directory tree are peripheral to this activity, and are used only as much as necessary to obtain access to files on behalf of a Windows client. The various cases the CIFS server can encounter while accessing files and directories, and what it does with UID/GID and permission bits in that access, are considered in the following sections.

Pre-existing directories and files

A pre-existing Linux directory will not have ACLs associated with it. In this case, the CIFS server will use the permission bits and the mapped UID/GID of the CIFS user to determine whether it has access to the directory contents. If the directory is written by the CIFS server, the inherited ACLS from the directory tree above that directory (if there are any) will be written into the directory so future CIFS access will have the ACLs to guide it.

Pre-existing files are treated like pre-existing directories. The CIFS server uses the UID/GID of the CIFS user and the permission bits to determine the access to the file. If the file is written to, the ACLs inherited from the containing directory for the file are applied to the file using the standard Windows ACL inheritance rules.

Working with pre-existing files and directories

Pre-existing file treatment has ramifications for cross-protocol environments. If, for example, files are deposited into a directory tree using NFS and then accessed using CIFS clients, the directory tree will not have ACLs associated with it, and access to the files will be moderated by the NFS UID/GID and permissions bits. If those files are then modified by a CIFS client, they will take on the UID/GID of the CIFS client (the new owner) and the NFS clients may lose access to those files.

New directories and files

New directories created in a tree by the Windows client inherit the ACLs of the parent directory. They ACLs are created with the UID/GID of the Windows user (the UID/GID that the SID for the Windows user is mapped to) and they have a Linux permission bit mask of 700. This translates to Linux applications (which do not understand the Windows ACL) having owner and group (users with the same group ID) with read, write, execute permissions, and everyone else having just read and execute permissions.

New files are handled the same way as directories. The files inherit the ACLs of the parent directory according to the Windows rules for ACL inheritance, and they are created with a UID/GID of the Windows user as mapped from the SID. They are assigned a permissions mask of 700.

Working with new files and directories

The inheritance rules of Windows assume that all directories are created on a Windows machine, where they inherit ACLs from their parent; the top level of a directory tree (the root of the file system) is assigned ACLs by the file system formatting process from the defaults for the system.

This process is not in place on file serving nodes. Instead, when the management console creates a share on a node, the share does not have any inherited ACLs from the root of the file system in which it is created. This leads to strange behavior when a Windows client attempts to use permissions to control access to a file in such a directory. The usual CREATOR/OWNER and EVERYBODY ACLs (which are a part of the typical Windows ACLS inheritance ACL set) do not exist on the containing directory for the share, and are not inherited downward into the share directory tree. For true Windows-like behavior, the creator of a share must access the root of the share and set the desired ACLs on it manually (using Windows Explorer or a command line tool such as ICACLS). This process is somewhat unnatural for Linux administrators, but should be fairly normal for Windows administrators. Generally, the administrator will need to create a CREATOR/OWNER ACL that is inheritable on the share directory, and then create an inheritable ACL that controls default access to the files in the directory tree.

Changing the way CIFS inherits permissions on files accessed from Linux applications

To avoid the CIFS server modifying file permissions on directory trees that a user wants to access from Linux applications (so keeping permissions other than 700 on a file in the directory tree), a user can set the `setgid` bit in the Linux permissions mask on the directory tree. When the `setgid` bit is set, the CIFS server honors that bit, and any new files in the directory inherit the parent directory permission bits and group that created the directory. This maintains group access for new files created in that directory tree until `setgid` is turned off in the tree. That is, Linux-style permissions semantics are kept on the files in that tree, allowing CIFS users to modify files in the directory while NFS users maintain their access though their normal group permissions.

For example, if a user wants all files in a particular tree to be accessible by a set of Linux users (say, through NFS), the user should set the `setgid` bit (through local Linux mechanisms) on the top level directory for a share (in addition to setting the desired group permissions, for example 770). Once that is done, new files in the directory will be accessible to the group that creates the directory and the permission bits on files in that directory tree will not be modified by the CIFS server. Files that existed in the directory before the `setgid` bit was set are not affected by the change in the containing directory; the user must manually set the group and permissions on files that already existed in the directory tree.

This capability can be used to facilitate cross-protocol sharing of files. Note that this does not affect the permissions inheritance and settings on the CIFS client side. Using this mechanism, a Windows user can set the files to be inaccessible to the CIFS users of the directory tree while opening them up to the Linux users of the directory tree.

X9000 Windows Clients and Windows ACLs

The X9000 Windows Client also stores and uses ACLs, but the format it uses is not compatible with the current CIFS ACL implementation. The X9000 Windows Client and the CIFS server operate in different permissions domains and do not typically honor each other's ACLs.

Troubleshooting CIFS

Changes to user permissions do not take effect immediately

The CIFS implementation maintains an authentication cache that is set to four hours. If a user is authenticated to a share, and the user's permissions are then changed, the old permissions will remain in effect until the cache expires, at four hours after the authentication. The next time the user is encountered, the new, correct value will be read and written to the cache for the next four hours.

This is not a common occurrence. However, to avoid the situation, use the following guidelines when changing user permissions:

- After a user is authenticated to a share, wait four hours before modifying the user's permissions.
- Conversely, it is safe to modify the permissions of a user who has not been authenticated in the previous four hours.

Robocopy errors occur during node failover or failback

If Robocopy is in use on a client while a file serving node is failed over or failed back, the application repeatedly retries to access the file and reports the error `The process cannot access the file because it is being used by another process`. These errors

occur for 15 to 20 minutes. The client's copy will then continue without error if the retry timeout has not expired. To work around this situation, take one of these steps:

- Stop and restart the Likewise process on the affected file serving node:

  ```
  # /opt/likewise/bin/lwsm stop lwreg && /etc/init.d/lwsmd stop
  # /etc/init.d/lwsmd start && /opt/likewise/bin/lwsm start srvsvc
  ```

- Power down the file serving node before failing it over, and do failback operations only during off hours.

The following `xcopy` and `robocopy` options are recommended for copying files from a client to a highly available CIFS server:

`xcopy`: include the option `/C`; in general, `/S /I /Y /C` are good baseline options.

`robocopy`: include the option `/ZB`; in general, `/S /E /COPYALL /ZB` are good baseline options.

Copy operations interrupted by node failback

If a node failback occurs while `xcopy` or `robocopy` is copying files to a CIFS share, the copy operation might be interrupted and need to be restarted.

Active Directory users cannot access CIFS shares

If any AD user is set to UID 0 in Active Directory, you will not be able to connect to CIFS shares and errors will be reported. Be sure to assign a UID other than 0 to your AD users.

8 Using FTP

The FTP feature allows you to create FTP file shares for data stored on the cluster. Clients access the FTP shares using standard FTP and FTPS protocol services.

IMPORTANT: Before configuring FTP, select an authentication method (either Local Users or Active Directory). See "Configuring authentication for CIFS, FTP, and HTTP" (page 42) for more information.

To configure FTP, first create one or more configuration profiles. A profile defines global FTP parameters and specifies the file serving nodes on which the parameters are applied.

You can then create FTP shares. A share defines parameters such as access permissions and lists the file system to be accessed through the share. Each share is associated with a configuration profile. The share parameters are added to the profile's global parameters on the file serving nodes specified in the configuration profile.

You can administer FTP from the management console GUI or the command line. On the GUI, select **FTP** from the Navigator to display the current FTP configuration. The Config Profile pane lists the profiles that have been created. The Share pane in the lower part of the screen shows the FTP shares associated with the selected profile.

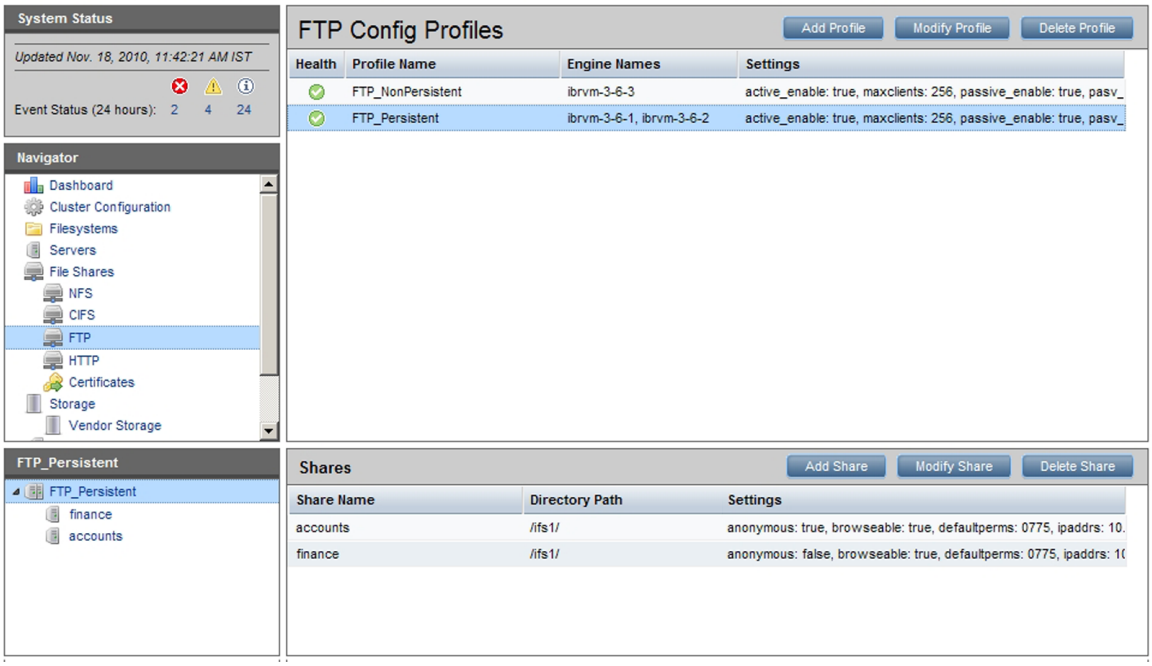

On the command line, FTP is managed by the `ibrix_ftpconfig` and `ibrix_ftpshare` commands. For more information, see the *HP StorageWorks X9000 File Serving Software CLI Reference Guide*.

Best practices for configuring FTP

When configuring FTP, follow these best practices:

- Create a configuration profile and then create the shares that will use that profile.
- If an SSL certificate will be required for FTPS access, add the SSL certificate to the cluster before creating the shares. See "Managing SSL certificates" (page 83) for information about creating certificates in the format required by X9000 Software and then adding them to the cluster.
- When configuring a share on a file system, the file system must be mounted.

- If the directory path to the share includes a subdirectory, be sure to create the subdirectory on the file system and assign read/write/execute permissions to it. (X9000 Software does not create the subdirectory if it does not exist, and instead adds a `/pub/` directory to the share path.)
- For High Availability, when specifying IP addresses for accessing a share, use IP addresses for VIFs having VIF backups. See the administrator guide for your system for information about creating VIFs.
- The allowed ports are 21 (FTP) and 990 (FTPS).

Managing configuration profiles

A configuration profile specifies a set of global FTP parameters that are in effect on the nodes listed in the profile. The `vsftpd` service starts on these nodes when the cluster services start. Only one configuration profile can be in effect on a particular node.

Adding a configuration profile

To add a configuration profile from the GUI, click **Add Profile** on the Config Profiles pane. On the Add FTP Profile dialog box, specify the appropriate parameters for the profile and select the servers where the profile will be active.

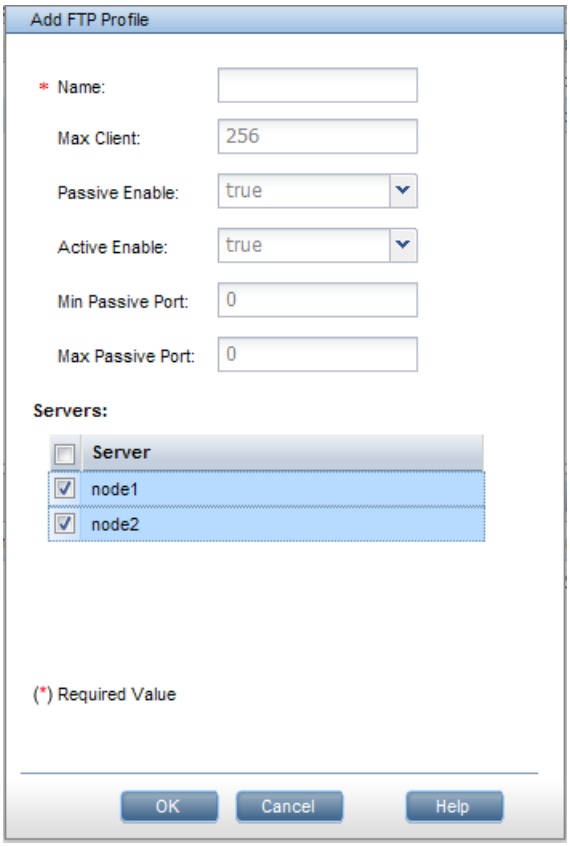

To add a configuration profile from the command line, use the following command:

```
ibrix_ftpconfig -a PROFILENAME [-h HOSTLIST] [-S SETTINGLIST]
```

For the `-S` option, use a comma to separate the settings, and enclose the settings in quotation marks, such as `"passive_enable=TRUE,maxclients=200"`. To see a list of available settings for the profile, use the following command:

```
ibrix_ftpconfig -L
```

Modifying a configuration profile

To modify a configuration profile, select the profile on the Config Profiles pane, and click **Modify Profile**. You can then make the necessary changes on the Modify FTP Profile dialog box. To modify a configuration profile from the command line, use the following command:

```
ibrix_ftpconfig -m PROFILENAME [-h HOSTLIST] [-S SETTINGLIST]
```

Viewing configuration profiles

The Modify FTP Profile dialog box shows details for a specific configuration profile. You can also use the following command to view detailed information about configuration profiles:

```
ibrix_ftpconfig -i -h HOSTLIST [-v level]
```

Deleting a configuration profile

To remove a configuration profile, select the profile on the FTP Config Profiles pane, click **Delete Profile**, and confirm the operation. From the command line, use the following command to delete a profile:

```
ibrix_ftpconfig -d PROFILENAME
```

Managing FTP shares

You can create an FTP share when you create a file system through the GUI. (See "Creating a file system" (page 11) for more information.) You can also add a share to an existing file system, as described here. You can create multiple shares having the same physical path, but with different sets of properties, and then assign users to the appropriate share. Be sure to use a different IP address or port for each share.

NOTE: The file system must be mounted when you create the share.

Each FTP share is associated with a configuration profile. The FTP properties associated with the share are added to the file serving nodes hosting the configuration profile.

Adding an FTP share

To add an FTP share from GUI, select the profile to be associated with the share, and then click **Add Share** on the Share pane.

On the General tab, specify the file system, which must be mounted, and the default directory path for the share. If the directory path includes a subdirectory, be sure to create the subdirectory on the file system and assign read/write/execute permissions to it. (X9000 Software does not create the subdirectory if it does not exist, and instead adds a `/pub/` directory to the share path.)

IMPORTANT: Ensure that all users who are given read or write access to HTTP shares have sufficient access permissions at the file system level for the directories exposed as shares.

Also select the appropriate FTP parameters for the share and enter the IP addresses and ports that clients will use to access the share. For High Availability, specify the IP address of a VIF having a VIF backup.

NOTE: The allowed ports are 21 (FTP) and 990 (FTPS).

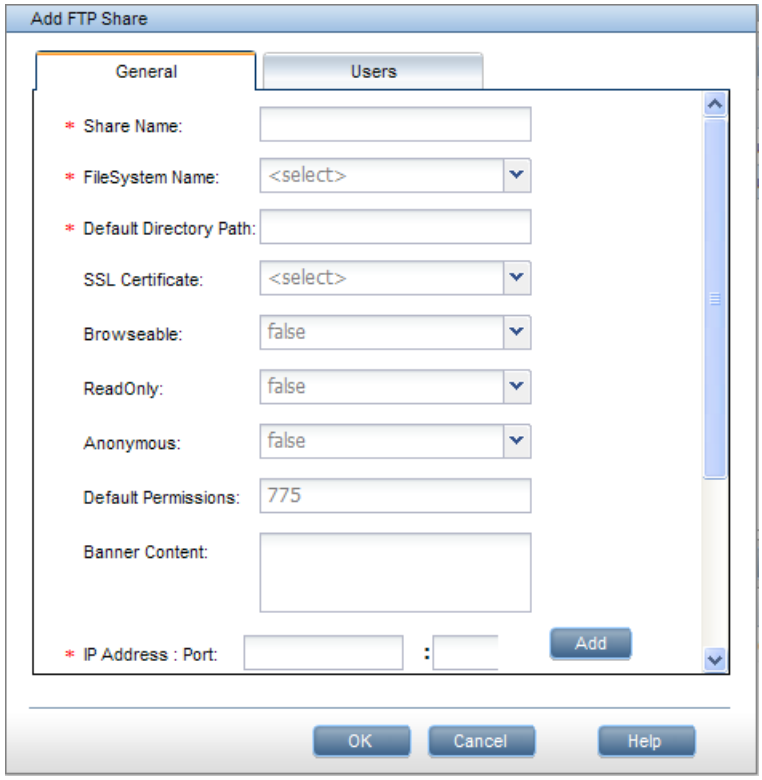

The Users tab lists the users allowed read access, write access, or both on a share directory. You can add or delete users as necessary.

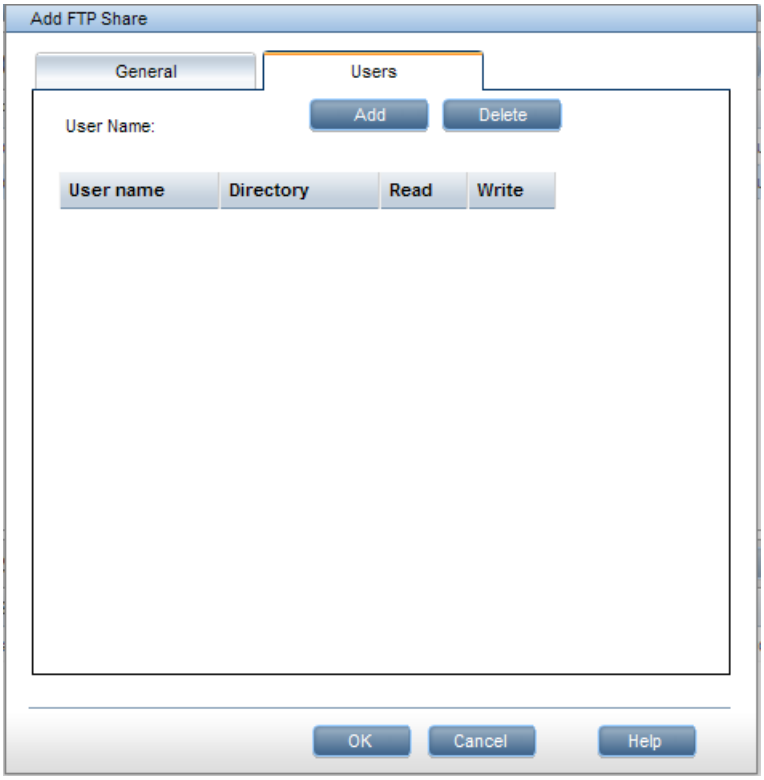

To assign share permissions to specific users, click **Add** and complete the Add Users to Share dialog box.

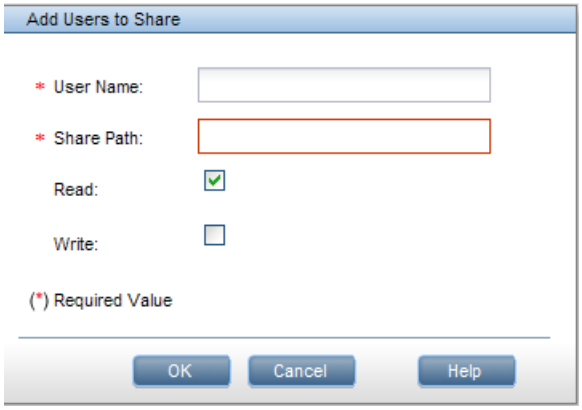

To add an FTP share from the command line, use the following command:

```
ibrix_ftpshare -a SHARENAME -c PROFILENAME -f FSNAME -p dirpath -I
IP-Address:Port [-u USERLIST] [-S SETTINGLIST]
```

For the `-S` option, use a comma to separate the settings, and enclose the settings in quotation marks, such as `"browseable=true,readonly=true"`. For the `-I` option, use a semicolon to separate the IP address:port settings and enclose the settings in quotation marks, such as `"ip1:port1;ip2:port2;..."`. To see a list of available settings for the share, use the following command:

```
ibrix_ftpshare -L
```

Modifying an FTP share

To change the properties for an FTP share, select the share on the Share pane, click **Modify Share**, and make the necessary changes on the Modify FTP Share dialog box. To modify a share from the CLI, use the following command:

```
ibrix_ftpshare -m SHARENAME -c PROFILENAME [-f FSNAME -p dirpath] -I
IP-Address:Port [-u USERLIST] [-S SETTINGLIST]
```

Viewing FTP shares

The Modify FTP Share dialog box shows the configuration of an FTP share. To view this information from the CLI, use the following command:

```
ibrix_ftpshare -i SHARENAME -c PROFILENAME  [-v level]
```

Deleting an FTP share

To remove an FTP share, select the share on the Share pane, click **Delete Share**, and confirm the operation. To remove the share from the command line, use the following command:

```
ibrix_ftpshare -d SHARENAME -c PROFILENAME
```

The vsftpd service

When the cluster services are started on a file serving node, the `vsftpd` service starts automatically if the node is included in a configuration profile. Similarly, when the cluster services are stopped, the `vsftpd` service also stops. If necessary, use the Linux command `ps -ef | grep vsftpd` to determine whether the service is running.

If you do not want `vsftpd` to run on a particular node, remove the node from the configuration profile.

IMPORTANT: For FTP share access to work properly, the `vsftpd` service must be started by X9000 Software. Ensure that the `chkconfig` of `vsftpd` is set to OFF (`chkconfig vsftpd off`).

Starting or stopping the FTP service manually

Use the following command to start the FTP service manually:

```
/usr/local/ibrix/ftpd/etc/vsftpd start /usr/local/ibrix/ftpd/hpconf/
```

Use the following command to stop the FTP service manually:

```
/usr/local/ibrix/ftpd/etc/vsftpd stop /usr/local/ibrix/ftpd/hpconf/
```

Use the following command to restart the FTP service manually:

```
/usr/local/ibrix/ftpd/etc/vsftpd restart /usr/local/ibrix/ftpd/hpconf/
```

NOTE: When the FTP configuration is changed with the management console GUI or CLI, the FTP daemon is restarted automatically.

Accessing shares

Clients can access an FTP share by specifying a URL in their browser (Internet Explorer or Mozilla Firefox). In the following URLs, `IP_address:port` is the IP (or virtual IP) and port configured for the share.

- For a share configured with an IP-based virtual host and the `anonymous` parameter is set to `true`, use the following URL:

  ```
  ftp://IP_address:port/
  ```

- For a share configured with a userlist and having the `anonymous` parameter set to `false`, use the following URL:

  ```
  ftp://<ADDomain\username>@IP_address:port/
  ```

NOTE: When a file is uploaded into an FTP share, the file is owned by the user who uploaded the file to the share.

If a user uploads a file to an FTP share and specifies a subdirectory that does not already exist, the subdirectory will not be created automatically. Instead, the user must explicitly use the `mkdir ftp` command to create the subdirectory. The permissions on the new directory are set to 777. If the anonymous user created the directory, it is owned by `ftp:ftp`. If a non-anonymous user created the directory, the directory is owned by `user:group`.

You can also use `curl` commands to access an FTP share. (The default SSL port is 990.)

For anonymous users:

- Upload a file using FTP protocol:

  ```
  curl -T <filename> -k ftp://IP_address/pub/ -u anonymous
  ```

- Upload a file using FTPS protocol:

  ```
  curl -T <filename> -k --ftp-ssl-reqd ftp://IP_address:990/pub/ -u ftp
  ```

- Download a file using FTP protocol:

  ```
  curl -k ftp://IP_address/pub/<filename> -u anonymous
  ```

- Download a file using FTPS protocol:

  ```
  curl -k --ftp-ssl-reqd ftp://IP_address:990/pub/<file_name> -u ftp
  ```

The following example shows an anonymous client accessing a share.

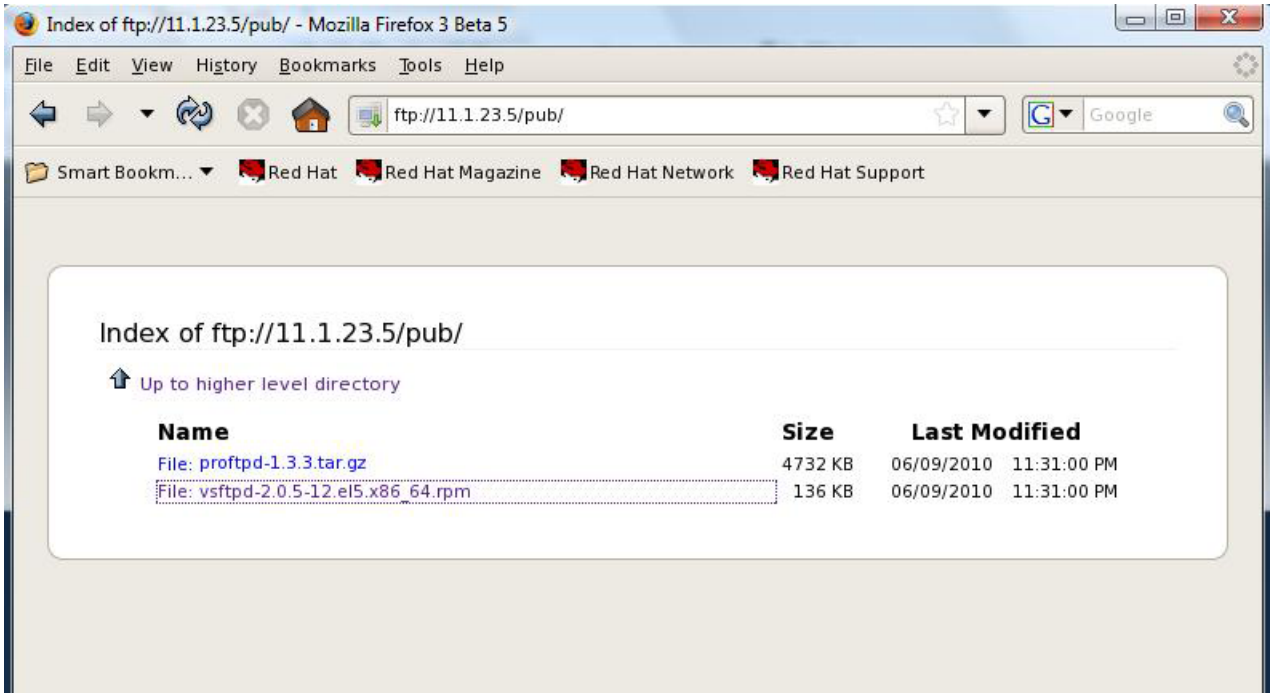

For Active Directory users (specify the user as in this example: `ASM2k3.com\\ib1`):

- Upload a file using FTP protocol:

  `curl -T <filename> -k ftp://IP_address/pub/ -u <ADuser>`

- Upload a file using FTPS protocol:

  `curl -T <filename> -k --ftp-ssl-reqd ftp://IP_address:990/pub/ -u <ADuser>`

- Download a file using FTP protocol:

  `curl -k ftp://IP_address/<filename> -u <ADuser>`

- Download a file using FTPS protocol:

  `curl -k --ftp-ssl-reqd ftp://IP_address:990/<filename> -u (ADuser>`

Shares can be accessed from any management console that has FTP clients:

`ftp <Virtual_IP>`

For FTPS, use the following command from the management console:

`lftp -u <user_name> -p <ssl port> -e 'set ftp:ssl-force true' <share_IP>`

9 Using HTTP

The HTTP feature allows you to create HTTP file shares for data stored on the cluster. Clients access the HTTP shares using standard HTTP and HTTPS protocol services.

IMPORTANT: Before configuring HTTP, select an authentication method (either Local Users or Active Directory). See "Configuring authentication for CIFS, FTP, and HTTP" (page 42) for more information.

The HTTP configuration consists of a configuration profile, a virtual host, and an HTTP share. A profile defines global HTTP parameters that apply to all shares associated with the profile. The virtual host identifies the IP addresses and ports that clients will use to access shares associated with the profile. A share defines parameters such as access permissions and lists the file system to be accessed through the share.

HTTP is administered from the management console GUI or the command line. On the GUI, select **HTTP** from the File Shares list in the Navigator. The Config Profile screen lists the current HTTP configuration, including the existing configuration profiles and the virtual hosts configured on the selected profile.

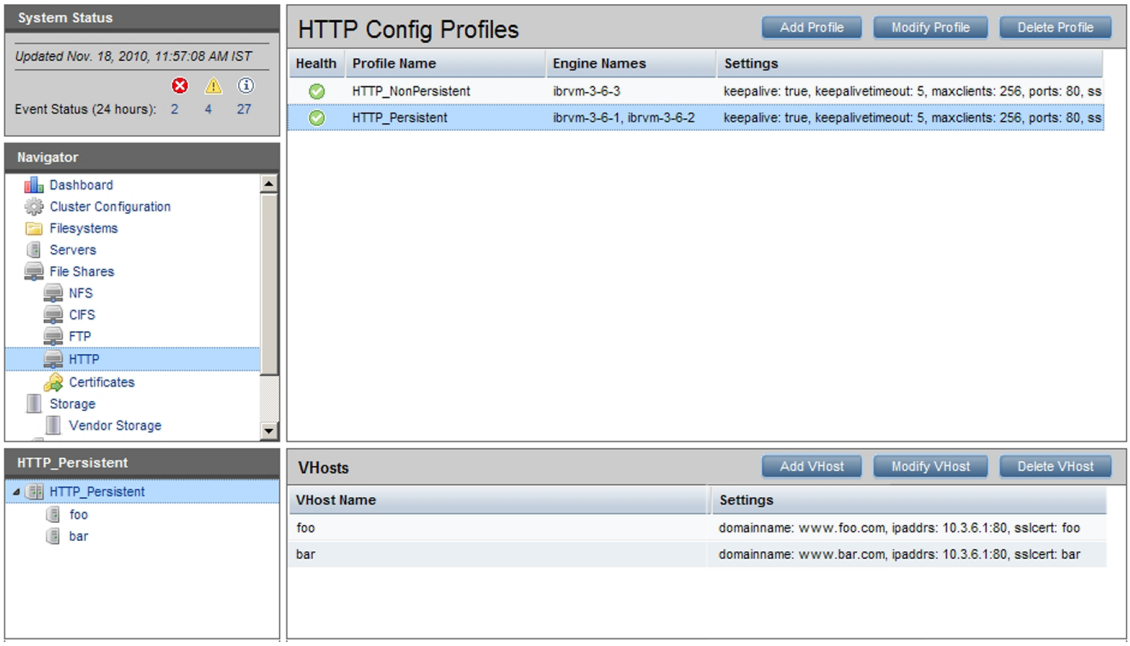

On the command line, HTTP is managed by the `ibrix_httpconfig`, `ibrix_httpvhost`, and `ibrix_httpshare` commands. For more information, see the *HP StorageWorks X9000 File Serving Software CLI Reference Guide*.

Best practices for configuring HTTP

When configuring HTTP, follow these best practices:

- If an SSL certificate will be required for HTTPS access, add the SSL certificate to the cluster before creating the shares. See "Managing SSL certificates" (page 83) for information about creating certificates in the format required by X9000 Software and then adding them to the cluster.
- When configuring a share on a file system, the file system must be mounted.
- If the directory path to the share includes a subdirectory, be sure to create the subdirectory on the file system and assign read/write/execute permissions to it. (X9000 Software does

not create the subdirectory if it does not exist, and instead adds a `/pub/` directory to the share path.)

- For High Availability, when specifying IP addresses for accessing a share, use IP addresses for VIFs having VIF backups. See the administrator guide for your system for information about creating VIFs.

Configuring HTTP with the HTTP Wizard

The New HTTP Wizard creates a configuration profile, a virtual host, and an HTTP share. To begin, click **Add Profile** on the Config Profiles pane. On the Create Profile dialog box, specify the appropriate parameters for the profile and select the servers to associate with the profile.

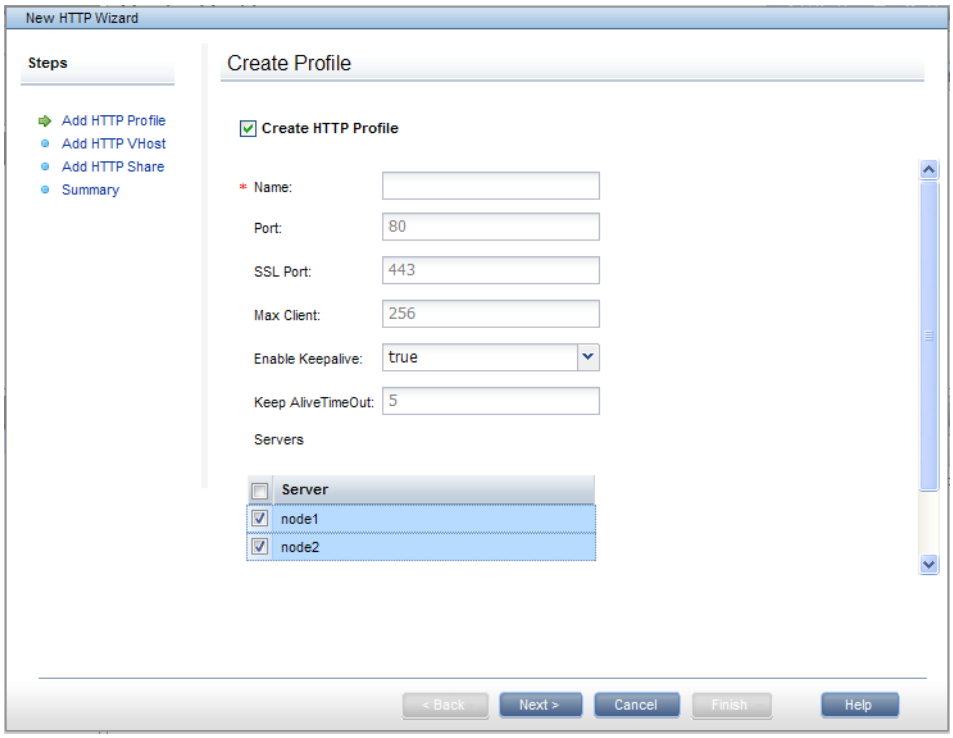

On the Create Vhost dialog box, click **Create HTTP Vhost** and then specify the appropriate parameters. You can add multiple IP addresses:ports on the user network for the virtual host. For High Availability, specify a VIF having a VIF backup.

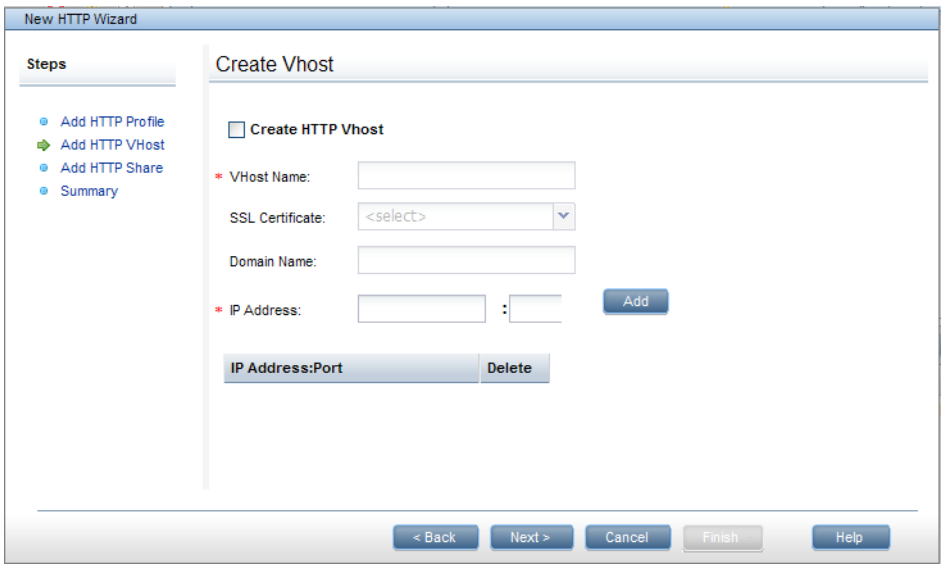

On the Create Share dialog box, click **Create HTTP Share** and set the appropriate parameters. Note the following:

- The file system selected for the share must be mounted. If the directory path includes a subdirectory, be sure to create the subdirectory on the file system and assign read/write/execute permissions to it. (X9000 Software does not create the subdirectory if it does not already exist, and instead adds a `/pub/` directory to the share path.)
- When specifying the **URL Path**, do not include `http://<IP address>` or any variation of this in the URL path. For example, `/reports/` is a valid URL path.
- Set the **Anonymous** field to false only if you want to restrict access to specific users.

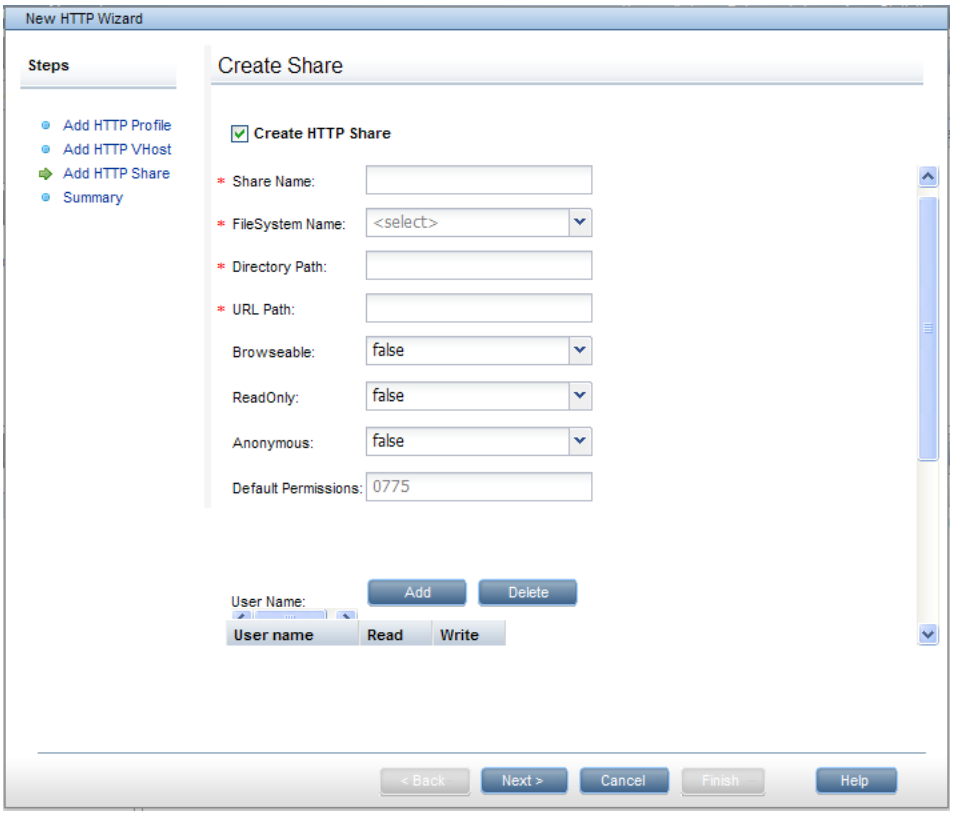

To allow specific users read access, write access, or both, click **Add**. On the Add Users to Share dialog, assign the appropriate permissions to the user. When you complete the dialog, the user is added to the list on the Create Share dialog box.

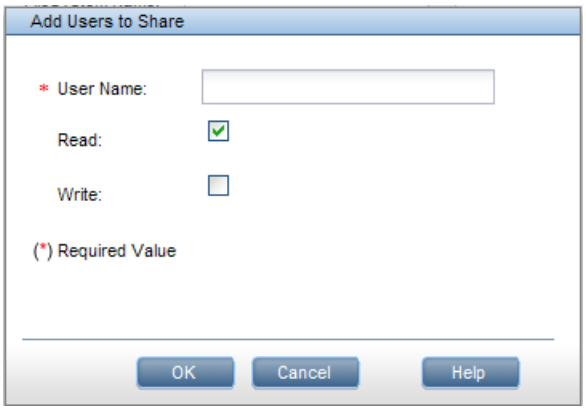

The Summary screen presents an overview of the HTTP configuration. You can go back and modify any part of the configuration if necessary.

When the wizard is complete, users can access the share from a browser. For example, if you configured the share with the anonymous user, specified 99.226.50.92 as the IP address on the Create Vhost dialog box, and specified `/reports/` as the URL path on the Create Share dialog box, users can access the share using the following URL:

```
http://99.226.50.92/reports/
```

The users will see an index of the share (if the browseable property of the share is set to true), and can open and save files. For more information about accessing shares and uploading files, see "Accessing shares" (page 81).

Managing configuration profiles

Adding a configuration profile from the CLI

To add a configuration profile from the command line, use the following command:

```
ibrix_httpconfig -a PROFILENAME [-h HOSTLIST] [-S SETTINGLIST]
```

For the `-S` option, use a comma to separate the settings, and enclose the settings in quotation marks, such as `"keepalive=true,maxclients=200,..."`. To see a list of available settings for the share, use `ibrix_httpconfig -L`.

Modifying a configuration profile

To modify a configuration profile, select the profile on the Config Profiles screen, click **Modify Profile**, and make the necessary changes on the Modify HTTP Profile dialog box. To modify a configuration profile from the command line, use the following command:

```
ibrix_httpconfig -m PROFILENAME [-h HOSTLIST] [-S SETTINGLIST]
```

Viewing configuration profiles

The Modify HTTP Profile dialog box shows details for a specific configuration profile. To view this information from the CLI, use the following command:

```
ibrix_httpconfig -i PROFILENAME [-v level]
```

Deleting a configuration profile

To remove a configuration profile, select the profile on the Config Profiles page, click **Delete Profile**, and confirm the operation. From the command line, use the following command to delete a profile:

```
ibrix_httpconfig -d PROFILENAME
```

Managing virtual hosts

If an SSL certificate is required for HTTP access, you can specify the certificate when you create the vhost. The SSL certificate must already exist. Add your certificates on the GUI (select **Certificates** from the Navigator), or use the `ibrix_certificate` command to add them.

Adding a virtual host

To add an HTTP virtual host, click **Add VHost** on the Vhost pane. The HTTP Wizard then opens at the Create Vhost dialog box. To add a virtual host from the CLI, use the following command:

```
ibrix_httpvhost -a VHOSTNAME -c PROFILENAME -I IP-Address:Port [-S
SETTINGLIST]
```

Modifying a virtual host

To change the properties for a virtual host, select the virtual host on the VHosts page, click **Modify Vhost**, and make the necessary changes on the Modify VHost Properties dialog box. To modify a virtual host from the CLI, use the following command:

```
ibrix_httpvhost -m VHOSTNAME -c PROFILENAME -I IP-Address:Port [-S SETTINGLIST]
```

Viewing a virtual host

The VHosts pane shows the configuration of virtual hosts. To view the configuration from the CLI, use the following command:

```
ibrix_httpvhost -i VHOSTNAME -c PROFILENAME [-v level]
```

Deleting a virtual host

To remove a virtual host, select the virtual host on the Vhosts pane, click **Delete Vhost**, and then confirm the operation. To remove the virtual host from the command line, use the following command:

```
ibrix_httpvhost -d VHOSTNAME -c PROFILENAME
```

Managing HTTP shares

You can create an HTTP share when you create a file system on the GUI (see "Creating a file system" for more information), or you can create a share with the HTTP Wizard. You can also add a share to an existing virtual host. For example, you might need to create multiple shares having the same physical path, but with different sets of properties, and then assign users to the appropriate share.

IMPORTANT: Ensure that all users who are given read or write access to HTTP shares have sufficient access permissions at the file system level for the directories exposed as shares.

NOTE: The file system must be mounted when you create the share.

To view HTTP shares on the GUI, select the appropriate profile on the Config Profile pane, and then select the appropriate virtual host from the lower navigator. The Share page shows the shares configured on that virtual host.

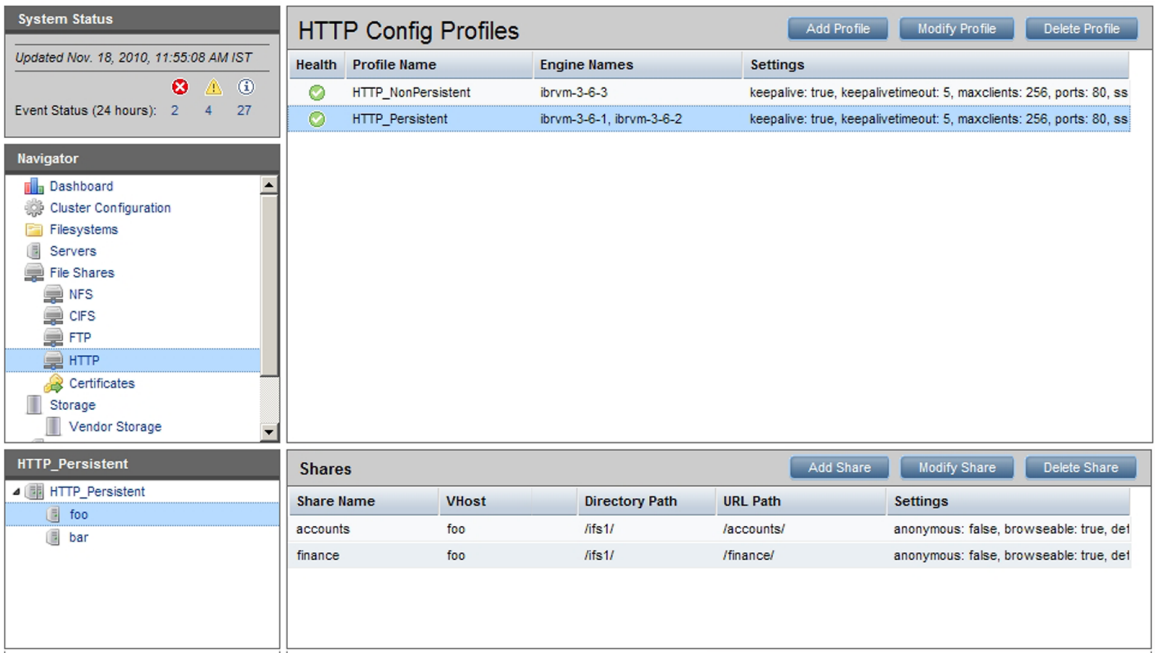

Adding an HTTP share

To add an HTTP share, click **Add Share** on the Share page. The HTTP Wizard opens at the Create Share dialog box.

To add an HTTP share from the command line, use the following command:

```
ibrix_httpshare -a SHARENAME -c PROFILENAME -t VHOSTNAME -f FSNAME -p
dirpath -P urlpath [-u USERLIST] [-S SETTINGLIST]
```

For the `-S` option, use a comma to separate the settings, and enclose the settings in quotation marks, such as `"browseable=true,readonly=true"`. To see the valid settings for an HTTP share, use the following command:

```
ibrix_httpshare –L
```

Modifying an HTTP share

To change the properties for an HTTP share, select the share on the Share page, click **Modify Share**, and make the necessary changes on the Modify HTTP Share dialog box. To modify a share from the CLI, use the following command:

```
ibrix_httpshare -m SHARENAME -c PROFILENAME -t VHOSTNAME [-f FSNAME -p
dirpath] [-P urlpath] [-u USERLIST] [-S SETTINGLIST]
```

Viewing HTTP shares

The Modify HTTP Share dialog shows the complete configuration of an HTTP share. To view this information from the CLI, use the following command:

```
ibrix_httpshare -i SHARENAME -c PROFILENAME -t VHOSTNAME [-v level]
```

Deleting an HTTP share

To remove an HTTP share, select the share on the Share page, click **Delete Share**, and then confirm the operation. To remove the share from the command line, use the following command:

```
ibrix_httpshare -d SHARENAME -c PROFILENAME -t VHOSTNAME
```

Starting or stopping the HTTP service manually

Use the following command to start the HTTP service manually:

```
/usr/local/ibrix/httpd/bin/apachectl -k start -f
/usr/local/ibrix/httpd/conf/httpd.conf
```

Use the following command to stop the HTTP service manually:

```
/usr/local/ibrix/httpd/bin/apachectl -k stop -f
/usr/local/ibrix/httpd/conf/httpd.conf
```

Use the following command to restart the HTTP service manually:

```
/usr/local/ibrix/httpd/bin/apachectl -k restart -f
/usr/local/ibrix/httpd/conf/httpd.conf
```

NOTE: When the HTTP configuration is changed with the management console GUI or CLI, the HTTP daemon is restarted automatically.

Accessing shares

Clients access an HTTP share by specifying a URL in their browser (Internet Explorer or Mozilla Firefox). In the following URLs, `IP_address:port` is the IP (or virtual IP) and port configured for the share.

- For a share configured with an IP-based virtual host and the `anonymous` parameter is set to `true`, use the following URL:

  ```
  http://IP_address:port/urlpath/
  ```

- For a shared configured with a `userlist` and having the `anonymous` parameter set to `false`, use the following URL:

  ```
  http://IP_address:port/urlpath/
  ```

  Enter your user name and password when prompted.

NOTE: When a file is uploaded into an HTTP share, the file is owned by the user who uploaded the file to the share.

If a user uploads a file to an HTTP share and specifies a subdirectory that does not already exist, the subdirectory will be created. For example, you could have a share mapped to the directory `/ifs/http/` and using the url `http_url`. A user could upload a file into the share:

```
curl -T file http://<ip>:<port>/http_url/new_dir/file
```

If the directory `new_dir` does not exist under `http_url`, the http service automatically creates the directory `/ifs/http/new_dir/` and sets the permissions to 777. If the anonymous user performed the upload, the `new_dir` directory is owned by `daemon:daemon`. If a non-anonymous user performed the upload, the `new_dir` directory is owned by `user:group`.

You can also use `curl` commands to access an HTTP share.

For anonymous users:

- Upload a file using HTTP protocol:

  ```
  curl -T <filename> http://IP_address:port/urlpath/
  ```

- Upload a file using HTTPS protocol:

  ```
  curl --cacert <cacert_file> -T <filename>
  https://IP_address:port/urlpath/<filename>/
  ```

- Download a file using HTTP protocol:

  ```
  curl http://IP_address:port/urlpath/<filename> -o <path to
  download>/<filename>/
  ```

- Download a file using HTTPS protocol:

  ```
  curl --cacert <cacert_file>
  https://IP_address:port/urlpath/<filename> -o <path to
  download>/<filename>/
  ```

For Active Directory users (specify the user as in this example: `mycompany.com\\User1`):

- Upload a file using HTTP protocol:

  ```
  curl –T <filename> -u <ADuser> http://IP_address:port/urlpath/
  ```

- Upload a file using HTTPS protocol:

  ```
  curl --cacert <cacert_file> -T <filename> -u <ADuser>
  https://IP_address:port/urlpath/
  ```

- Download a file using HTTP protocol:

  ```
  curl -u <ADuser> http://IP_address/dils/urlpath -o path to
  download>/<filename>/
  ```

- Download a file using HTTPS protocol:

  ```
  curl --cacert <cacert_file> -u <ADuser>
  https://IP_address:port/urlpath/<file_name> -o path to
  download>/<filename>/
  ```

10 Managing SSL certificates

Servers accepting FTPS and HTTPS connections typically provide an SSL certificate that verifies the identity and owner of the web site being accessed. You can add your existing certificates to the cluster, enabling file serving nodes to present the appropriate certificate to FTPS and HTTPS clients. X9000 Software supports PEM certificates.

When you configure the FTP share or the HTTP vhost, select the appropriate certificate.

You can manage certificates from the GUI or the CLI, On the GUI, select **Certificates** from the Navigator to open the Certificates panel. The Certificate Summary shows the parameters for the selected certificate.

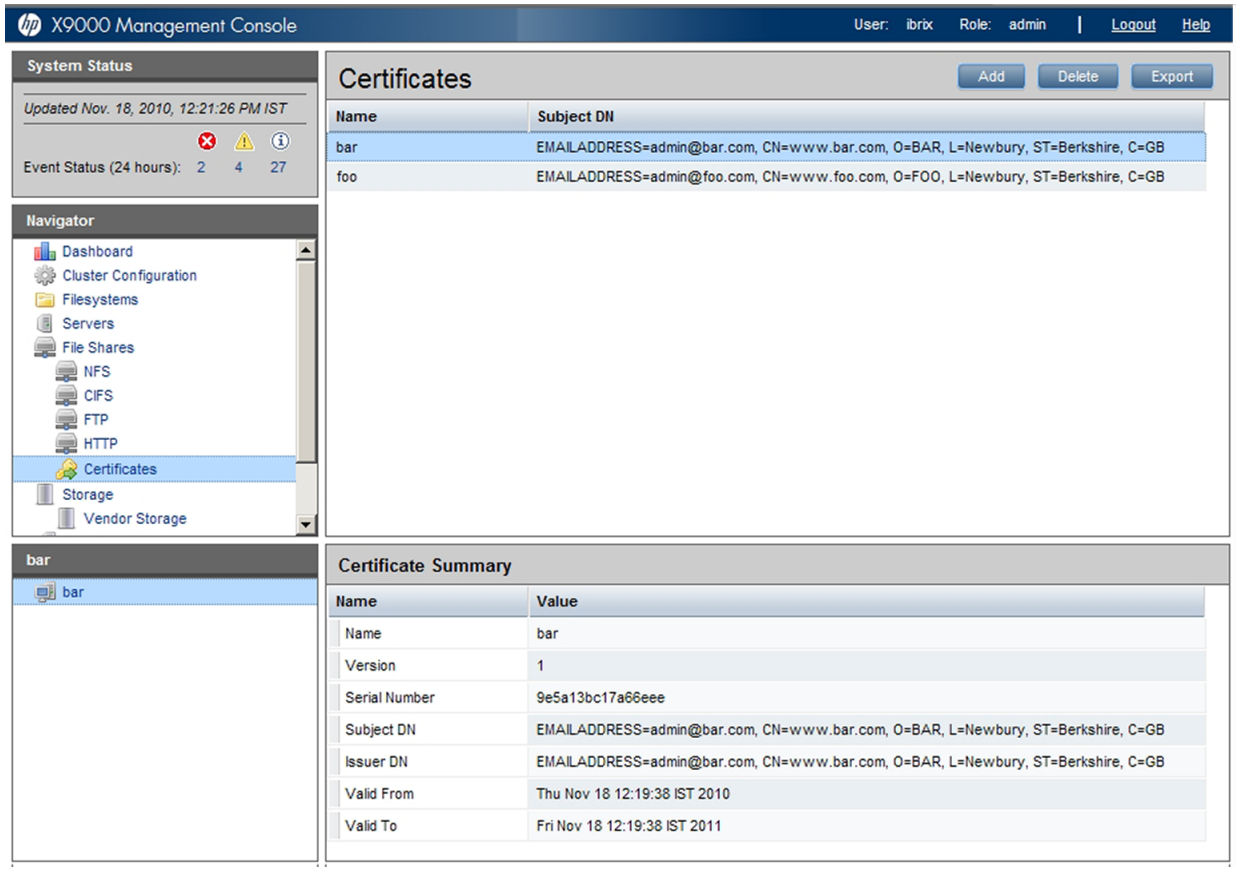

Creating an SSL certificate

Before creating a certificate, OpenSSL must be installed and must be included in your PATH variable (in RHEL5, the path is `/usr/bin/openssl`).

There are two parts to a certificate: the certificate contents (specified in a `.crt` file) and a private key (specified in a `.key` file). Certificates added to the cluster must meet these requirements:

- The certificate contents (the `.crt` file) and the private key (the `.key` file) must be concatenated into a single file.
- The concatenated certificate file must include the headers and footers from the `.crt` and `.key` files.
- The concatenated certificate file cannot contain any extra spaces.

Before creating a real certificate, you can create a self-signed SSL certificate and test access with it. Complete the following steps to create a test certificate that meets the requirements for use in an X9000 cluster:

1. Generate a private key:

   ```
   openssl genrsa -des3 -out server.key 1024
   ```

   You will be prompted to enter a passphrase. Be sure to remember the passphrase.

2. Remove the passphrase from the private key file (`server.key`). When you are prompted for a passphrase, enter the passphrase you specified in step 1.

   ```
   cp server.key server.key.org
   openssl rsa -in server.key.org -out server.key
   rm -f server.key.org
   ```

3. Generate a Certificate Signing Request (CSR):

   ```
   openssl req -new -key server.key -out server.csr
   ```

4. Self-sign the CSR:

   ```
   openssl x509 -req -days 365 -in server.csr -signkey server.key -out
   server.crt
   ```

5. Concatenate the signed certificate and the private key:

   ```
   cat server.crt server.key > server.pem
   ```

When adding a certificate to the cluster, use the concatenated file (`server.pem` in our example) as the input for the GUI or CLI.

The following example shows a valid PEM encoded certificate that includes the certificate contents, the private key, and the headers and footers:

```
-----BEGIN CERTIFICATE-----
MIICUTCCAboCCQCIHW1FwFn2ADANBgkqhkiG9w0BAQUFADBtMQswCQYDVQQGEwJV
UzESMBAGA1UECBMJQmVya3NoaXJlMRAwDgYDVQQHEwdOZXdidXJ5MQwwCgYDVQQK
EwNhYmMxDDAKBgNVBAMTA2FiYzEcMBoGCSqGSIb3DQEJARYNYWRtaW5AYWJjLmNv
bTAeFw0xMDEyMTEwNDQ0MDdaFw0xMTEyMTEwNDQ0MDdaMG0xCzAJBgNVBAYTAlVT
MRIwEAYDVQQIEwlCZXJrc2hpcmUxEDAOBgNVBAcTB05ld2J1cnkxDDAKBgNVBAoT
A2FiYzEMMAoGA1UEAxMDYWJjMRwwGgYJKoZIhvcNAQkBFg1hZG1pbkBhYmMuY29t
MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQDdrjHH/W93X7afTIUOrllCHw21
u31tinMDBZzi+R18r9SZ/muuyvG4kJCbOoQnohuir/s4aAEULAOnf4mvqLfZlkBe
25HgT+ImshLzyHqPImuxTEXvjG5H1sEDLNuQkHvl8hF9Wxao1tv4eL8TL5KqK1W6
8juMVAw2cFDHxji2GQIDAQABMA0GCSqGSIb3DQEBBQUAA4GBAKvYJK8RXKMObCKk
ae6oJ36FEkdl/ACHCw0Nxk/VMR4dv9lIk8Dv8sdYUUqHkNAME2yOaRI190c5bWSa
MjhSjOOqUmmgmeDYlAu+ps3/1Fte5yl4ZV8VCu7bHCWx2OSy46Po03MMOu99JXrB
/GCKE8fO8Fhyq/7LjFDR5GeghmSw
-----END CERTIFICATE-----
-----BEGIN RSA PRIVATE KEY-----
MIICXgIBAAKBgQDdrjHH/W93X7afTIUOrllCHw21u31tinMDBZzi+R18r9SZ/muu
yvG4kJCbOoQnohuir/s4aAEULAOnf4mvqLfZlkBe25HgT+ImshLzyHqPImuxTEXv
jG5H1sEDLNuQkHvl8hF9Wxao1tv4eL8TL5KqK1W68juMVAw2cFDHxji2GQIDAQAB
AoGBAMXPWryKeZyb2+np7hFbompOK32vAA1vLZHUwFoI0Tch7yQ60vv2PBvlZCQf
4y06ik5xmkqLA+tsGxarx8DnXKUy0PHJ3hu6mTocIJdqqN0n+KO4tG2dvDPdSE7l
phX2sY9MVt4X/QN3eNb/F3cHjnM9BYEr0BY3mTkKXz61jzABAkEA+M3PProYwvS6
P8m4DenZh6ehsu4u/ycjmW/ujdp/PcRd5HBAWJasTXTezF5msugHnnNBe8F1i1q4
9PfL0C+kuQJBAOQXjrmPZxDc8YA/V45MUKv4eHHN0E03p84budtblHQ70BCLaO41
n267t3DrZfW+VtsVDVBMja4UhoBasgv3rGECQQCILDR6k2YMBd+OG/xleRD6ww+o
G96S/bvpNa7t6qFrj/cHmTxOgCDLv+RVHHG/B2lsGo7Dig2oeL30LU9aoUjZAkBV
KSqDw7PyitusS3oQShQQsTufGf385pvDi3yQFxhNcYuUschisCivumyaP3mZEBDz
yV9oLLz1UvqI79PsPfPhAkEAxSqebd1Ymqr2wi0RnKTmHfDCb3yWLPi57kc+lgrK
LUlxawhTzDwzTWJ9m4gQqRlAaXoIElfk6ITwW0g9Th5Ouw==
-----END RSA PRIVATE KEY-----
```

NOTE: When you are ready to create a real SSL certificate, consult the following site for a description of the procedure:

http://httpd.apache.org/docs/2.0/ssl/ssl_faq.html#selfcer

Adding a certificate to the cluster

To add an existing certificate to the cluster, click **Add** on the Certificates panel. On the Add Certificate dialog box, enter a name for the certificate. Use a Linux command such as `cat` to display your concatenated certificate file. For example:

```
cat server.pem
```

Copy the contents of the file to the Certificate Content section of the dialog box. The copied text must include the certificate contents and the private key in PEM encoding. It must also include the proper headers and footers, and cannot contain any extra spaces.

NOTE: You can add only one certificate at a time.

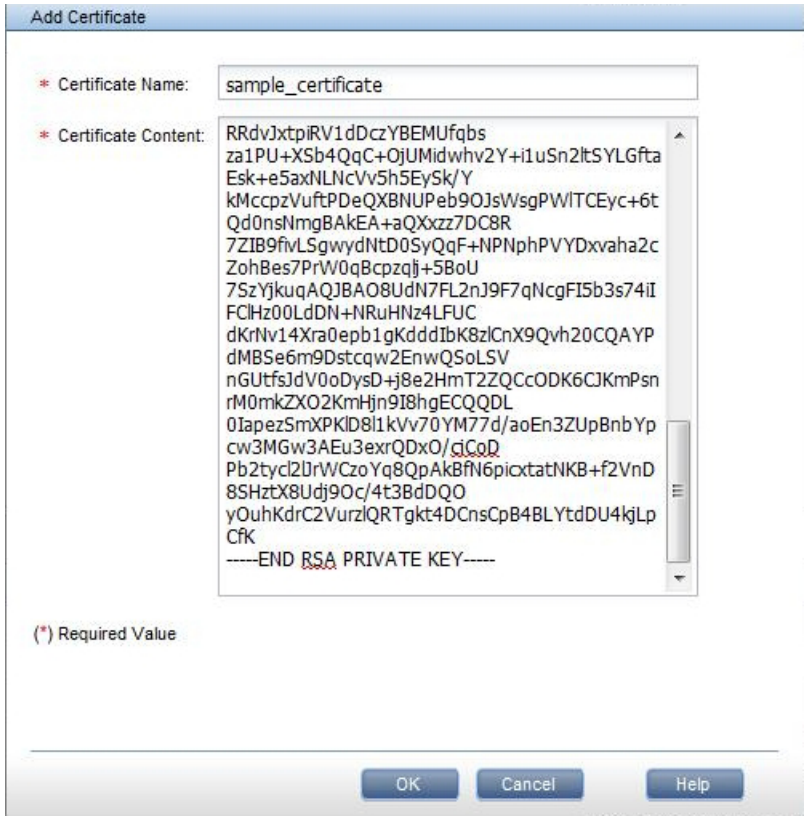

The certificate is saved on all file serving nodes in the directory `/usr/local/ibrix/pki`. If your cluster uses a dedicated Management Server to host the management console, the certificate is saved in `/usr/local/ibrix/certificates/certificates.xml` on that machine.

To add a certificate from the CLI, use the following command.

```
ibrix_certificate -a -c CERTNAME -p CERTPATH
```

For example:

```
# ibrix_certificate -a -c mycert -p server.pem
```

Run the command from the active management console. To add a certificate for a different node, copy that certificate to the active management console and then add it to the cluster. For example, if node `ib87` is hosting the active management console and you have generated a certificate for node `ib86`, copy the certificate to `ib87`:

```
scp server.pem ib87/tmp
```

Then, on node `ib87`, add the certificate to the cluster:

```
ibrix_certificate -a -c cert86 -p /tmp/server.pem
```

Exporting a certificate

If necessary, you can display a certificate and then copy and save the contents for future use. This step is called *exporting*. Select the certificate on the Certificates panel and click **Export**.

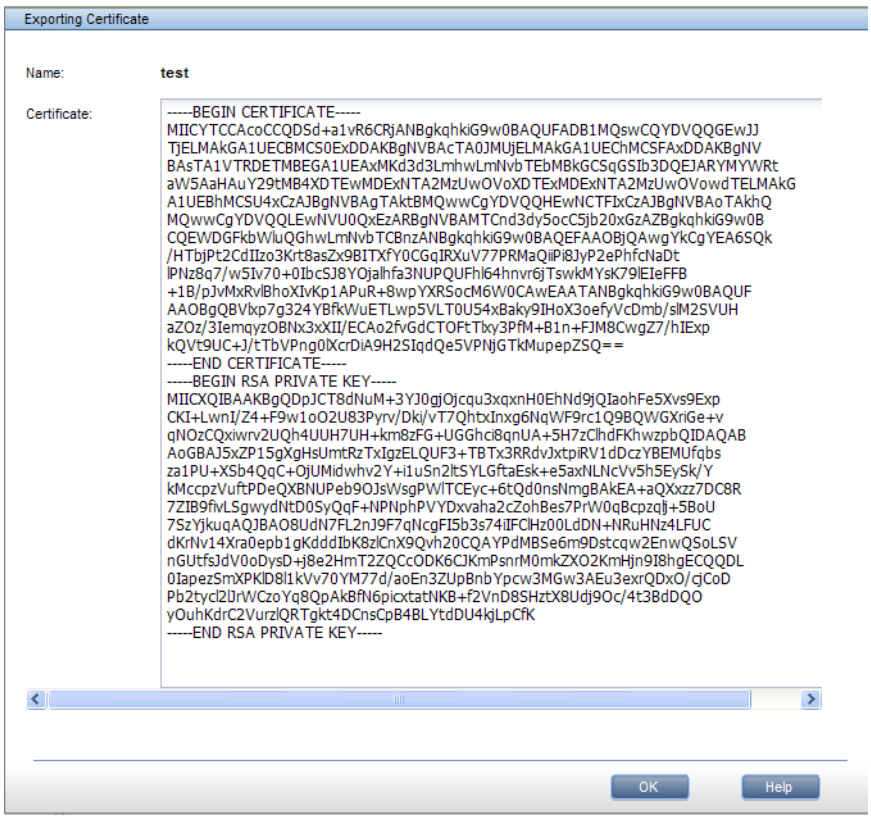

To export a certificate from the CLI, use this command:

```
ibrix_certificate -e -c CERTNAME
```

Deleting a certificate

To delete a certificate from the GUI, select the certificate on the Certificates panel, click **Delete**, and confirm the operation.

To delete a certificate from the CLI, use this command:

```
ibrix_certificate -d -c CERTNAME
```

11 Using remote replication

This chapter describes how to configure and manage remote replication.

Overview

Remote replication provides a transparent method to replicate changes in a source file system on one cluster to a target file system on either the same cluster or a second cluster. The remote replication service has two modes: continuous and run-once.

Both files and directories can be replicated with remote replication, and no special configuration of segments is needed. A remote replication task includes the initial synchronization of the source and target file systems.

When selecting file systems for remote replication, you should be aware of the following:

- One, multiple, or all file systems in a single cluster can be replicated.
- Remote replication is a one-way process. Bidirectional replication of a single file system is not supported.
- The mountpoint of the source file system can be different from the mountpoint on the target file system.

Remote replication has minimal impact on normal cluster operations:

- Cluster expansion (adding a new server) is allowed as usual on both the source and target.
- File systems can be exported over NFS, CIFS, FTP, or HTTP.
- Source or target file systems can be rebalanced while a remote replication job is in progress.
- File system policies (`ibrix_fs_tune`) can be set on both the source and target without any restrictions.

The management console initializes remote replication. However, each file serving node runs its own replication and synchronization processes, independent of and in parallel with other file serving nodes. The `ibrcfrd` daemons running on the file serving nodes perform the actual file system replication.

The source-side management console monitors the replication and watches for errors, failures, and so on.

Continuous or run-once replication modes

Remote replication can be used in two modes: continuous or run-once.

Continuous remote replication. This method tracks changes on the source file system and continuously replicates these changes to the target file system. The changes are tracked for the entire file system and are replicated in parallel by each file serving node. There is no strict order to replication at either the file-system or segment level. The continuous remote replication program tries to replicate on a first-in, first-out basis.

When you configure continuous remote replication, you must specify a file system as the source. You can specify either a file system or a file system and directory as the target. The following requirements apply:

- File systems specified as the replication source or target must exist.
- If a directory is specified as the replication target, the directory must exist on the target.

Run-once remote replication. This method allows for replication of a single directory sub-tree or an entire file system from the source file system to the target file system. Run-once is a point-in-time replication of all files and subdirectories within the specified directory or file system. All changes

that have occurred since the last replication job are replicated from the source file system to the target file system.

The following requirements apply to run-once replication:

- File systems specified as the replication source or target must exist.
- If a directory is specified as the replication target, the directory must exist on the target.
- When you specify a directory as the replication source, the actual replication location on the target is created by appending the entire path of the source directory to the specified target path. For example, if you specify the following:
  replication source: `<source_fs_root>dir1/dir2`
  replication target: `<target_fs_root>dir4/dir5`

  The contents of `<source_fs_root>dir1/dir2` is replicated to `<target_fs_root>dir4/dir5/dir1/dir2/`.
- If a directory is specified as the replication source, the directory path can have multiple levels of subdirectories and the full path must exist on the source. Using the previous example, `<source_fs_root>dir1/dir2` must exist on the source and `<target_fs_root>dir4/dir5` must exist on the target, but `dir1` or `dir1/dir2` does not need to exist on the target.

Remote cluster or intracluster

A replication task can use one of the following targets:

- **Remote cluster.** Configure either continuous or run-once replication.
- **The same cluster and a different file system.** Configure either continuous or run-once replication.
- **The same cluster and the same file system.** Configure run-once replication.

When configuring run-once replication for the same cluster and the same file system, be sure to specify two different, non-overlapping subdirectories as the source and target. For example, the following replication is *not* allowed:

From `<fs_root>dir1` to `<fs_root>dir1/dir2`

However, the following replication *is* allowed:

From `<fs_root>dir1` to `<fs_root>dir3/dir4`

Many-to-many or many-to-one replications

You can set up multiple replication tasks at once, each with a source/target replication pair. The tasks can be many-to-many or many-to-one replications.

Many-to-many replications

In many-to-many replications, each file system on the source cluster is replicated to a different file system on the remote target cluster. Note the following:

- Several replication tasks can run in parallel.
- Each replication is either a continuous or a run-once task.

Many-to-one replications

Many-to-one replications can be configured as follows:

- From multiple source clusters to different file systems on the same remote target cluster.
- From multiple source clusters to different subdirectories of the same target cluster and file system.

Note the following:

- Several replication tasks can run in parallel.
- Each replication is either a continuous or a run-once task.
- The target subdirectories must not overlap. For example, if source cluster 1 is replicating to `<fs_root>dir1`, source cluster 2 can replicate to `<fs_root>dir2`, but it cannot replicate to `<fs_root>dir1/dir2`.

Configuring the target for remote replication

You can configure remote replication on the management console GUI or from the CLI. Use the following steps to configure a target (a file system or directory) for remote replication. In this procedure, *target file system* refers to either a file system or a directory.

NOTE: The steps are not required when configuring intracluster replication.

- **Register source and destination clusters.** A cluster must be registered before other remote replication commands can use it as a target.
- **Export the target file system.** This step identifies the target file system for replication and associates it with the source cluster. Before replication can take place, you must create a mapping between the source cluster and the target file system that receives the replicated data. This mapping ensures that only the specified source cluster can write to the target file system.
- **Identify the preferred hosts for replication storage and a preferred NIC for replication traffic (optional).** Select the hosts and network NIC, or use the default preferences. The default preferences are as follows:
  - Use all hosts that have the file system mounted.
  - Use the cluster NIC on each host.

GUI procedure

This procedure must be run from the target cluster, and is not required for intracluster replication.

Select the file system on the GUI, and then select **Remote Replication Exports** from the lower Navigator. On the Remote Replication Exports tab, click **Export**. The Create Remote Replication Export dialog box allows you to specify the target file system for the replication.

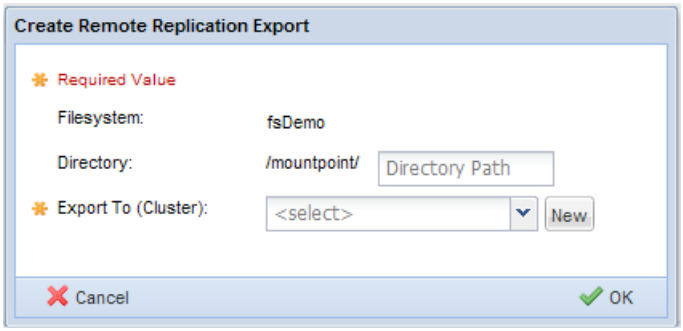

For **Export To (Cluster)**, select the source cluster of the file system to be replicated. If the source cluster is not in the list, you will need to register the cluster. Click **New** to open the Add Remote Cluster dialog box, and then specify the source cluster name and the IP address of the management console for that cluster.

NOTE: If the cluster uses an agile management console configuration, specify the virtual cluster name and the IP address of the virtual interface.

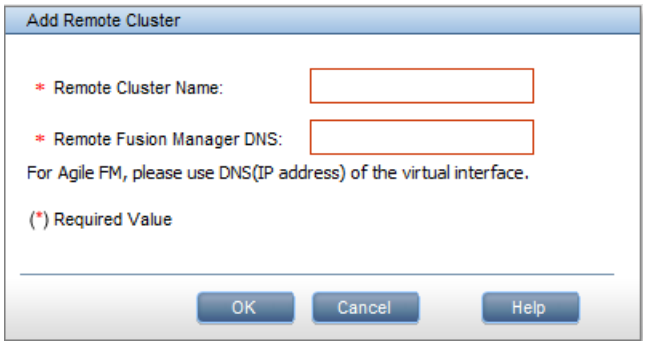

The GUI configures the replication export with the default target hosts and network interface preferences. To specify your own preferences, use the `ibrix_exportcfrpreference` command. For more information, see "Identifying host and NIC preferences on the target cluster" (page 91).

CLI procedure

NOTE: This procedure does not apply to intracluster replication.

Use the following commands to configure the target file system for remote replication:

1. Register the source and target clusters with the `ibrix_cluster -r` command if needed. To list the known remote clusters, run `ibrix_cluster -l` on the source cluster.
2. Export the target file system. Identify the target file system for the replication and associate it with the source cluster using the `ibrix_exportcfr` command.
3. Identify host and NIC preferences on the target cluster using the `ibrix_exportcfrpreference` command.

Registering source and target clusters

To register the clusters manually, use the following command:

`ibrix_cluster -r -C` *`CLUSTERNAME`* `-H` *`REMOTE_FM_HOST`*

`CLUSTERNAME` is the name of the management console for a cluster.

If the cluster is using an agile management console configuration, specify the `clusterName` displayed by the `ibrix_fm_tune -l` command, and enter the IP address of the cluster VIF.

NOTE: For each remote replication pair, cluster A must be registered with cluster B, and cluster B must be separately registered with cluster A.

To deregister a remote replication cluster, use the following command:

`ibrix_cluster –d –C` *`CLUSTERNAME`*

To list clusters registered on this management console, use the following command:

`ibrix_cluster -l`

Exporting the target file system

To create a mapping between the source cluster and the target file system where the replicated data will be stored, execute the following command on the target cluster:

`ibrix_exportcfr –f` *`FSNAME`* `[-p` *`DIRECTORY`*`] –C` *`REMOTE_CLUSTER`* `[–P]`

`FSNAME` is the target file system to be exported. The `–p` option can be used to export a directory located under the root of the specified file system. The `-C` *`REMOTE_CLUSTER`* option specifies the

source cluster (`clusterName`) containing the file system to be replicated. Include the `-P` option if you do not want to set the host and NIC preferences.

To unexport a file system for remote replication, use the following command:

```
ibrix_exportcfr –U –f TARGET_FSNAME
```

To list the remote replication exports, use the following command:

```
ibrix_exportcfr -l
```

Identifying host and NIC preferences on the target cluster

To identify the preferred hosts for replication storage and, optionally, a preferred NIC for replication traffic, use the following command:

```
ibrix_exportcfrpreference -a -f FSNAME -h HOSTLIST [-n IBRIX_NIC]
```

When specifying resources, note the following:

- Specify hosts by their host name or IP address. A host is any host on the target cluster that has the target file system mounted.
- Specify the network using the X9000 Software network name. Enter a valid user NIC or the cluster NIC. Setting the network preference is optional. If it is not specified, the host name (or IP) is used to determine the network.

The listed hosts will receive remote replication data via the specified NIC. To increase capacity, you can expand the number of preferred hosts by executing this command again with another list of hosts.

You can also use the `ibrix_exportcfrpreference` command for the following tasks:

- Restore the default preferences for remote replication:

  ```
  ibrix_exportcfrpreference -D -f FSNAME
  ```

- View preferences for remote replication. The output lists the exported file systems and associated host preferences on this cluster. All preferred hosts/NICs are listed with a corresponding ID number that can be used in commands to remove preferences.

  ```
  ibrix_exportcfrpreference -l
  ```

- Remove a remote replication preference:

  ```
  ibrix_exportcfrpreference -r -p PREFERENCE_ID
  ```

  To obtain the ID for a particular preferred host, list the remote replication preferences using the `-l` option.

Configuring and managing replication tasks on the GUI

NOTE: When configuring replication tasks, be sure to following the guidelines described in "Overview" (page 87).

To view or control replication tasks for a particular file system, select that file system on the GUI and then click **Tasks** > **Remote Replication** in the lower Navigator. The Remote Replication Tasks tab lists any replication tasks currently running or paused on the file system.

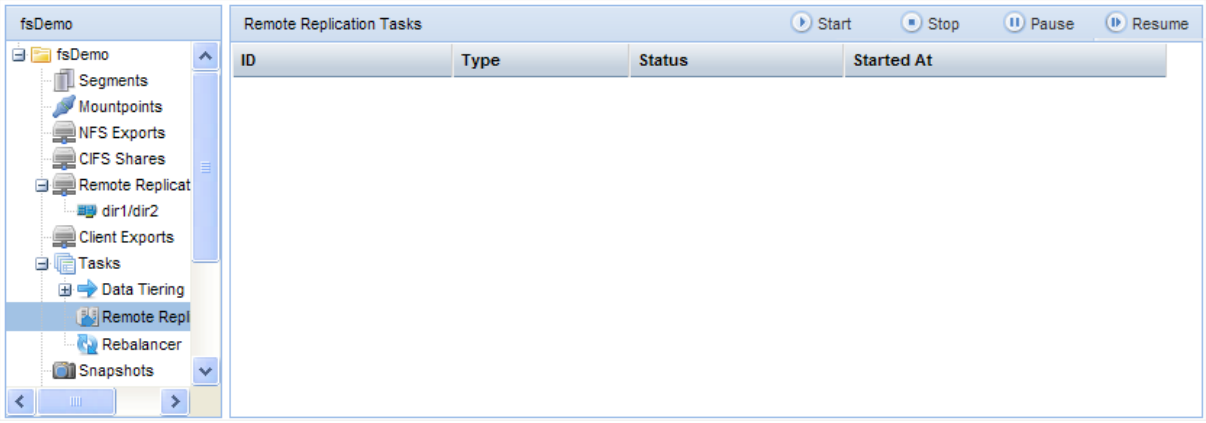

To stop a task, select the task and click **Stop**.

To pause a task, select the task and click **Pause**. The status of the task will change to PAUSED. Pausing a task that involves continuous data capture does not stop the data capture. You must allocate space on the disk to avoid running out of space because the data is captured but not moved. To resume a paused replication task, select the task and click **Resume**. The status of the task will change to RUNNING, and the task will continue from the point where it was paused.

To start a replication task, click **Start** to open the Start Remote Replication dialog box. Select the target for replication (a remote cluster, the same cluster, or the same cluster and file system), specify whether this is a continuous or run-once replication, and specify the target-side information. If the target cluster is not registered with the source cluster, click **New** and then register the cluster on the Add Remote Cluster dialog box.

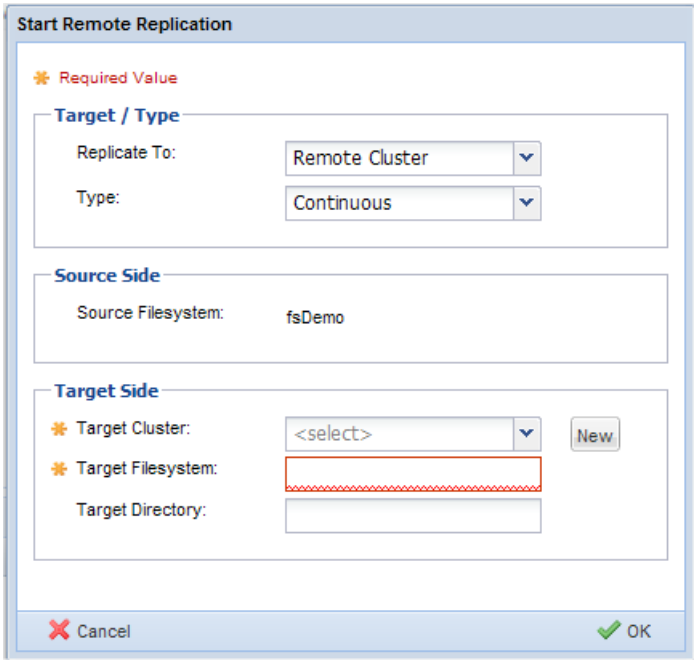

When the replication task begins, it will appear on the Remote Replication Tasks tab.

Configuring and managing replication tasks from the CLI

NOTE: When configuring replication tasks, be sure to following the guidelines described in "Overview" (page 87).

Starting a remote replication task to a remote cluster

Use the following command to start remote replication on the specified source file system. The command is executed from the source cluster.

```
ibrix_cfrjob -s -f SRC_FSNAME [-o [-S SRCDIR] ] -C TGT_CLUSTERNAME -F TGT_FSNAME
[-P TGTDIR]
```

The `-f` option specifies the source file system to be replicated. The `-C` option specifies the target cluster. If you are replicating to a directory on the target, `-P` specifies the target directory. If the `-P` option is not used, the mount point of the target filesystem will be used as the root of the replicated data.

Use the `-o` option for run-once jobs. This option can be used to synchronize single directories or entire file systems on the source and target in a single pass. If you do not specify a source directory with the `-S` option, the replication starts at the root of the file system. The run-once job terminates after the replication is complete; however, the job can be stopped manually, if necessary.

Starting an intra-cluster remote replication task

Use the following command to start a continuous or run-once intra-cluster replication task:

```
ibrix_cfrjob -s -f SRC_FSNAME [-o [-S SRCDIR]] -F TGT_FSNAME [-P TGTDIR]
```

The command starts a continuous or run-once intra-cluster replication task for file system *SRC_FSNAME*. The `-F` option specifies the name of the target file system (the default is the same as the source file system). The `-P` option specifies the target directory under the target file system (the default is the root of the file system).

Use the `-o` option to start a run-once job. The `-S` option specifies a directory under the source file system to synchronize with the target directory.

Starting a run-once directory replication task

Use the following command to start a run-once directory replication for file system *SRC_FSNAME*. The `-S` option specifies the directory under the source file system to synchronize with the target directory. The `-P` option specifies the target directory.

```
ibrix_cfrjob -s -f SRC_FSNAME -o -S SRCDIR -P TGTDIR
```

Stopping a remote replication task

Use the following command to stop a continuous or run-once replication task. Use the `ibrix_task -l` command to obtain the appropriate ID.

```
ibrix_cfrjob -k –n TASKID
```

Pausing a remote replication task

USe the following command to pause a continuous replication or run-once replication task with the specified task ID. Use the `ibrix_task -l` command to obtain the appropriate ID.

```
ibrix_cfrjob -p –n TASKID
```

Resuming a remote replication task

Use the following command to resume a continuous or run-once replication task with the specified task ID. Use the `ibrix_task -l` command to obtain the appropriate ID.

```
ibrix_cfrjob -r –n TASKID
```

Querying remote replication tasks

Use the following command to list all running and stopped continuous replication jobs in the cluster, optionally restricted by the specified file system and host name.

```
ibrix_cfrjob -l [-f SRC_FSNAME] [-h HOSTNAME] [-C SRC_CLUSTERNAME]
```

To see more detailed information about continuous and run-once replication tasks, run `ibrix_cfrjob` with the `-i` option. The command can optionally be restricted to a specified file system and host name. The display shows the status of jobs on each node, as well as job summary statistics (number of files in the queue, number of files processed). The query also indicates whether scanning is in progress on a given file serving node and lists any error conditions.

```
ibrix_cfrjob -i [-f SRC_FSNAME] [-h HOSTNAME] [-C SRC_CLUSTERNAME]
```

The following command prints detailed information about continuous replication tasks matching the specified task IDs. Use the `-h` option to limit the output to the specified host.

```
ibrix_cfrjob -i –n TASKIDS [ [-h HOSTNAME] [-C SRC_CLUSTERNAME]
```

Viewing replication status and activity

To view replication status and statistics, execute the `ibrix_cfrjob -i` command. If the command is executed from the source side, that management console answers the query with information gathered from the `ibrcfrd` daemons on the source-side file serving nodes.

Configuring remote failover/failback

When remote replication is configured from a local cluster to a remote cluster, you can fail over the local cluster to the remote cluster:

1. Stop write traffic to the local site.
2. Wait for all remote replication queues to drain.
3. Stop remote replication on the local site.
4. Reconfigure shares as necessary on the remote site. The cluster name and IP addresses (or VIFs) are different on the remote site, and changes are needed to allow clients to continue to access shares.
5. Redirect write traffic to the remote site.

When the local cluster is healthy again, take the following steps to perform a failback from the remote site:

1. Stop write traffic to the remote site.
2. Set up Run-Once remote replication, with the remote site acting as the source and the local site acting as the destination.
3. When the Run-Once replication is complete, restore shares to their original configuration on the local site, and verify that clients can access the shares.
4. Redirect write traffic to the local site.

Troubleshooting remote replication

Continuous remote replication fails when a private network is used

Continuous remote replication will fail if the configured cluster interface and the corresponding cluster Virtual Interface (VIF) for the management console are in a private network on either the source or target cluster. By default continuous remote replication uses the cluster interface and the Cluster VIF (the `ibrixinit –C` and `–v` options, respectively) for communication between the source cluster and the target cluster. To work around potential continuous remote replication communication errors, it is important that the `ibrixinit -C` and `-v` arguments correspond to a public interface and a public cluster VIF, respectively. If necessary, the

`ibrix_exportcfrpreference` command can be used to change the network interface preference.

12 Creating snapshots

The snapshot feature allows you to capture a point-in-time copy of a file system for online backup purposes and to simplify recovery of files from accidental deletion. The snapshot replicates all file system entities at the time of capture and is managed exactly like any other file system.

The snapshot feature is supported as follows:

- HP StorageWorks X9320 Network Storage System: supported on the HP StorageWorks P2000 G3 MSA Array System or HP StorageWorks 2000 Modular Smart Array G2 provided with the platform.
- HP StorageWorks X9300 Network Storage Gateway: supported on the HP StorageWorks P2000 G3 MSA Array System, HP StorageWorks 2000 Modular Smart Array G2, and Dell EqualLogic storage array (no arrays are provided with the X9300 system).
- HP StorageWorks X9720 Network Storage System: no support.

The snapshot feature uses the copy-on-write method to preserve the snapshot regardless of changes to the origin file system. Initially, the snapshot points to all blocks that the origin file system is using (**B** in the following diagram). When a block in the origin file system is overwritten with additions, edits, or deletions, the original block (prior to changes) is copied to the *snapshot store*, and the snapshot points to the copied block (**C** in the following diagram). The snapshot continues to point to the origin file system contents from the point in time that the snapshot was executed.

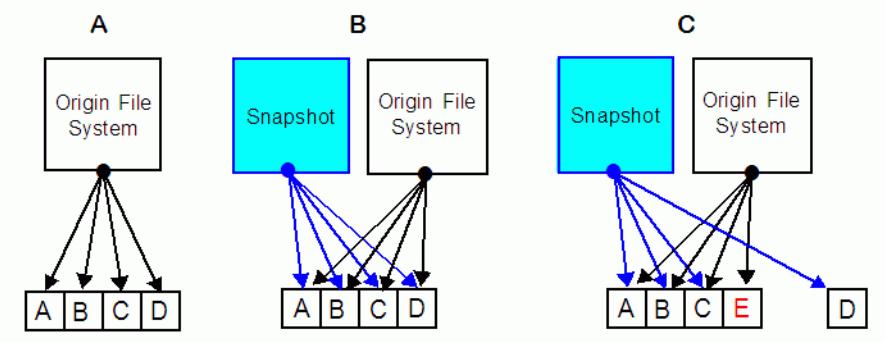

To create a snapshot, first provision or register the snapshot store. You can then create a snapshot from type-specific storage resources. The snapshot is active from the moment it is created.

You can take snapshots via the X9000 Software scheduler or manually, whenever necessary. Each snapshot maintains its origin file system contents until deleted from the system.

Snapshots can be made visible to users, allowing them to access and restore files (based on permissions) from the available snapshots.

NOTE: By default, snapshots are read only. HP recommends that you do not allow writes to any snapshots.

Setting up snapshots

This section describes how to configure the cluster to take snapshots.

Preparing the snapshot partition

The snapshot feature does not require any custom settings for the partition. However, HP recommends that you provide sufficient storage capacity to support the snapshot partition.

NOTE: If the snapshot store is too small, the snapshot will eventually exceed the available space (unless you detect this and manually increase storage). If this situation occurs, the array software deletes the snapshot resources and the X9000 Software snapshot feature invalidates the snapshot file system.

Although you can monitor the snapshot and manually increase the snapshot store as needed, the safest policy is to initially provision enough space to last for the expected lifetime of the snapshot. The optimum size of the snapshot store depends on usage patterns in the origin file system and the length of time you expect the snapshot to be active. Typically, a period of trial and error is required to determine the optimum size.

See the array documentation for procedures regarding partitioning and allocating storage for file system snapshots.

Registering for snapshots

After setting up the snapshot partition, you can register the partition with the management console. You will need to provide a name for the storage location and specify access parameters (IP address, user name, and password).

The following command registers and names the array's snapshot partition on the management console. The partition is then recognized as a repository for snapshots.

```
<installdirectory>/bin/ibrix_vs -r -n STORAGENAME -t { msa | eqlogic} -I IP(s) -U USERNAME
[-P PASSWORD]
```

To remove the registration information from the configuration database, use the following command. The partition will then no longer be recognized as a repository for snapshots.

```
<installdirectory>/bin/ibrix_vs -d -n STORAGENAME
```

Discovering LUNs in the array

After the array is registered, use the `-a` option to map the physical storage elements in the array to the logical representations used by X9000 Software. The software can then manage the movement of data blocks to the appropriate snapshot locations on the array.

Use the following command to map the storage information for the specified array:

```
<installdirectory>/bin/ibrix_vs -a [-n STORAGENAME]
```

Reviewing snapshot storage allocation

Use the following command to list all of the array storage that is registered for snapshot use:

```
<installdirectory>/bin/ibrix_vs -l
```

To see detailed information for named snapshot partitions on either a specific array or all arrays, use the following command:

```
<installdirectory>/bin/ibrix_vs -i [-n STORAGENAME]
```

Automated snapshots

If you plan to take a snapshot of a file system on a regular basis, you can automate the snapshots. To do this, first define an automated snapshot scheme, and then apply the scheme to the file system and create a schedule.

A snapshot scheme specifies the number of snapshots to keep and the number of snapshots to mount. You can create a snapshot scheme from either the management console GUI or the CLI.

The type of storage array determines the maximum number of snapshots you can keep and mount per file system.

Array	Maximum number of snapshots to keep	Maximum number of snapshots to mount
P2000 G3 MSA System/MSA2000 G2 array	32 snapshots per file system	7 snapshots per file system
EqualLogic array	8 snapshots per file system	7 snapshots per file system

For the P2000 G3 MSA System/MSA2000, the storage array itself also limits the total number of snapshots that can be stored. Arrays count the number of LUNs involved in each snapshot. For example, if a file system has four LUNs, taking two snapshots of the file system increases the total snapshot LUN count by eight. If a new snapshot will cause the snapshot LUN count limit to be exceeded, an error will be reported, even though the file system limits may not be reached. The snapshot LUN count limit on P2000 G3 MSA System/MSA2000 arrays is 255.

Creating automated snapshots using the GUI

On the Filesystems page of the management console GUI, select the file system where the snapshots will be taken and then, in the Navigator, select **Snapshots**. When the Snapshots page appears, click **Create**.

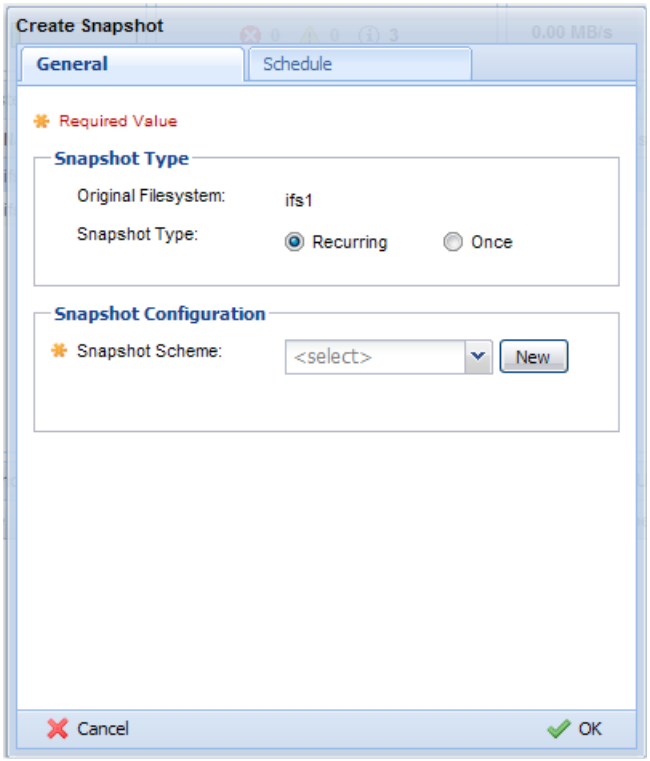

On the General tab, select **Recurring** as the Snapshot Type. Then, click **New** to create a new snapshot scheme. The Create Snapshot Scheme dialog box appears.

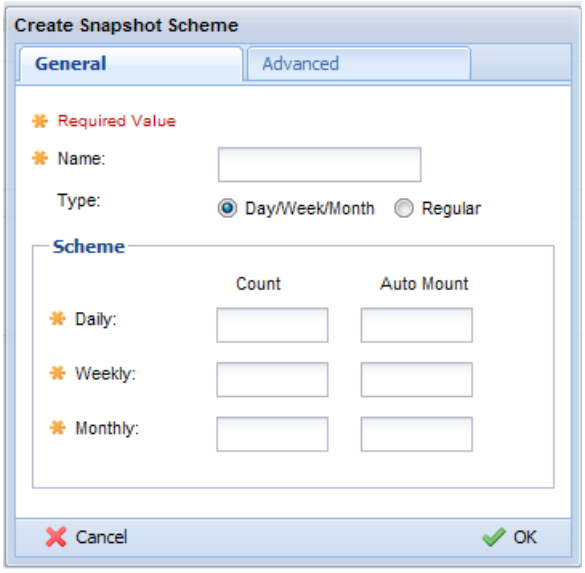

On the General tab, enter a name for the strategy. For the Type, select either **Day/Week/Month** or **Regular**.

- For **Day/Week/Month**, specify the number of snapshots to keep and mount on a daily, weekly, and monthly basis. Keep in mind the maximums allowed for your array type.

  NOTE: *Daily* means that one snapshot is kept per day for the specified number of days. For example, if you enter 6 as the daily count, the snapshot feature keeps 1 snapshot per day through the 6th day. On the 7th day, the oldest snapshot is deleted. Similarly, *Weekly* specifies the number of weeks that snapshots are retained, and *Month* specifies the number of months that snapshots are retained.

- For **Regular**, specify the total number of snapshots to keep and mount, keeping in mind the maximums allowed for your array type.

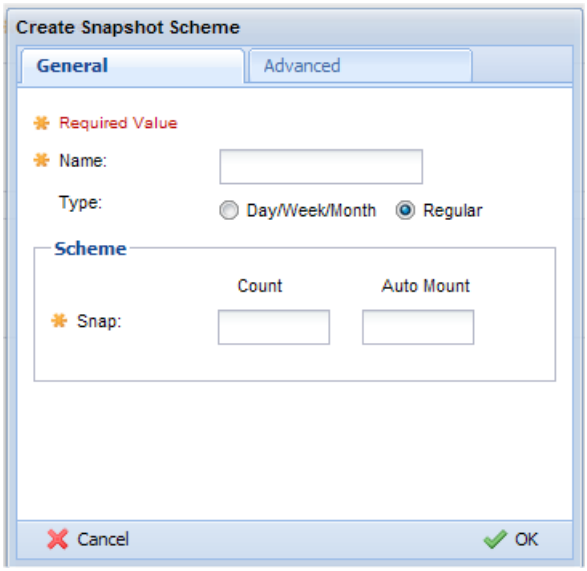

On the Advanced tab, you create templates for naming the snapshots and mountpoints. This step is optional.

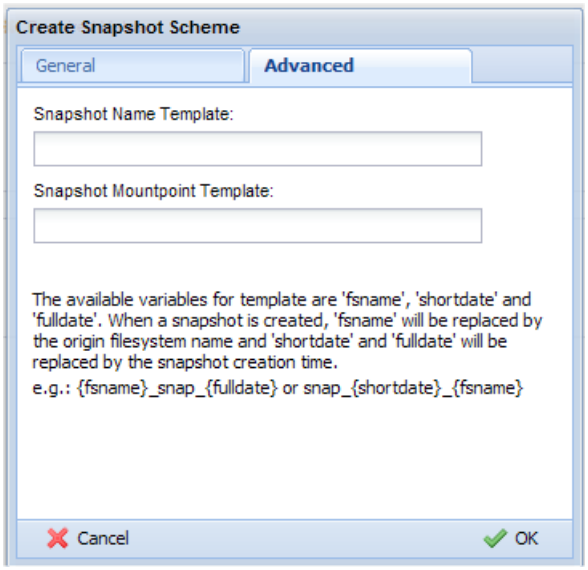
Create Snapshot Scheme
General
Advanced
Snapshot Name Template:
Snapshot Mountpoint Template:
The available variables for template are 'fsname', 'shortdate' and 'fulldate'. When a snapshot is created, 'fsname' will be replaced by the origin filesystem name and 'shortdate' and 'fulldate' will be replaced by the snapshot creation time.
e.g.: {fsname}_snap_{fulldate} or snap_{shortdate}_{fsname}
Cancel
OK

For either template, enter one or more of the following variables. The variables must be enclosed in braces ({ }) and separated by underscores (_). The template can also include text strings. When a snapshot is created using the templates, the variables are replaced with the following values.

Variable	Value
`fsname`	File system name
`shortdate`	*yyyy_mm_dd*
`fulldate`	*yyyy_mm_dd_HHmmz* + GMT

After creating the snapshot strategy, you can create a schedule for the automated snapshots. The new snapshot strategy will appear in the list of snapshot schemes on the General tab of the Create Snapshot dialog box. When you select the strategy, a description of the strategy will be displayed.

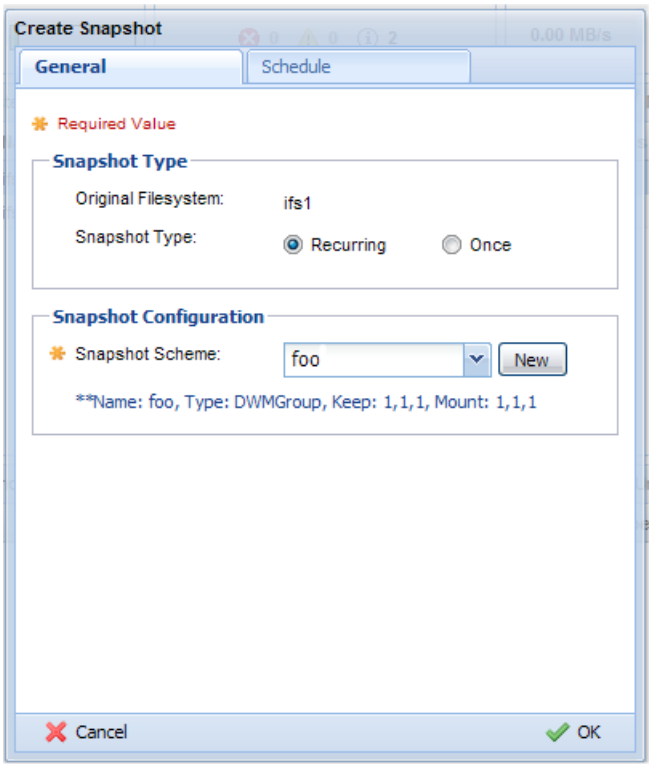

To create the schedule for the automated snapshots, go to the Schedule tab and click **Schedule this task (or run now)**. You can set the frequency of the snapshots and schedule when they should occur. You can also set start and end dates for the schedule. When you click **OK**, the snapshot scheduler will begin taking snapshots according to the specified snapshot strategy and schedule.

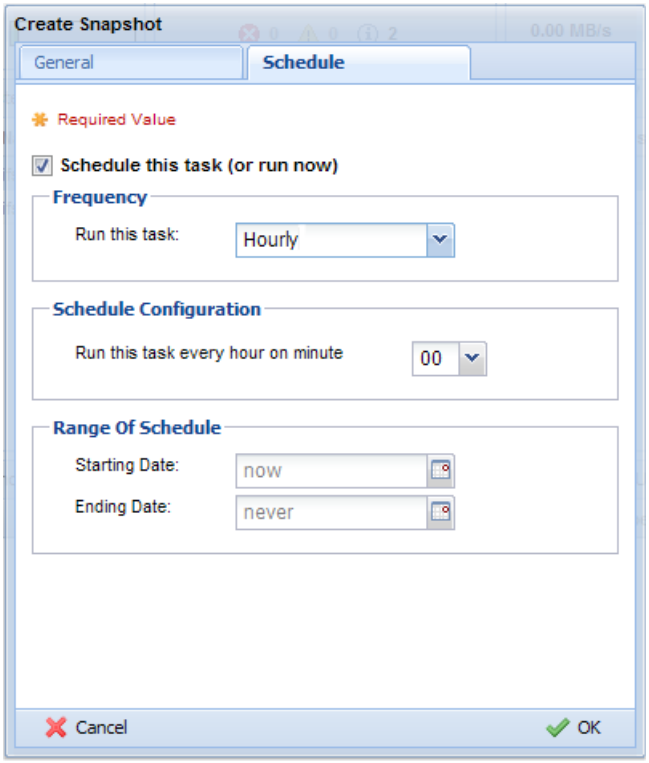

Creating automated snapshots from the CLI

You will first need to create an automated snapshot scheme with the `ibrix_snap_strategy` command and then schedule the snapshots with `ibrix_at`.

Creating a snapshot scheme

Use the following information to create an automated snapshot scheme from the CLI.

To define a snapshot scheme, execute the `ibrix_snap_strategy` command with the `-c` option:

```
ibrix_snap_strategy -c -n NAME -k KEEP -m MOUNT [-T TYPE] [-N NAMESPEC] [-M MOUNTSPEC]
```

The options are:

`-n` *NAME*	The name for the snapshot scheme.
`-T` *TYPE*	The strategy type: `linear` (corresponds to Regular on the GUI) or `DWMGroup`. `linear` allows a specific number of snapshots and mountpoints on the system. `DWMGroup` allows a specific number of snapshots and mountpoints per day, week, and month.

`-k KEEP`	The number of snapshots to keep per file system. For the P2000 G3 MSA System/MSA2000 G2 array, the maximum is 32 snapshots per file system. For Dell EqualLogic arrays, the maximum is eight snapshots per file system. If the strategy type is `linear`, enter the number of snapshots to keep, such as `-k 6`. If the strategy type is `DWMGroup`, enter the number of days, weeks, and months to retain snapshots. The numbers must be separated by commas, such as `-k 2,7,28`. **NOTE:** When the `DWMGroup` strategy is used, one snapshot is kept per day for the specified number of days. For example, if you enter 6 as the daily count, the snapshot feature keeps 1 snapshot per day through the 6th day. On the 7th day, the oldest snapshot is deleted. Similarly, the weekly count specifies the number of weeks that snapshots are retained, and the monthly count specifies the number of months that snapshots are retained.
`-m MOUNT`	The number of snapshots to mount per file system. The maximum number of snapshots is 7 per file system. If the strategy type is `linear`, enter the number of snapshots to mount, such as `-m 7`. If the strategy type is `DWMGroup`, enter the number of snapshots to mount per day, week, and month. The numbers must be separated by commas, such as `-m 2,2,3`. The sum of the numbers must be less than or equal to 7.
`-N NAMESPEC`	Snapshot name template. The template specifies a scheme for creating unique names for the snapshots. Use the variables listed below for the template.
`–M MOUNTSPEC`	Snapshot mountpoint template. The template specifies a scheme for creating unique mountpoints for the snapshots. Use the variables listed below for the template.

Variables for snapshot name and mountpoint templates.

`fulldate`	*yyyy_mm_dd_HHmmz* + GMT
`shortdate`	*yyyy_mm_dd*
`type`	Strategy type specified by the `-T` option (`linear` or `DWMGroup`)
`strategy`	Same as `type` but with the name of the “snapshot strategy” appended
`fsname`	File system name

You can specify one or more of these variables, enclosed in braces ({ }) and separated by underscores (_). The template can also include text strings. Two sample templates follow. When a snapshot is created using one of these templates, the variables will be replaced with the values listed above.

```
{fsname}_snap_{fulldate}
snap_{shortdate}_{fsname}
```

Scheduling and starting automated snapshots

The `ibrix_at` command is used to create the snapshot schedule. The command has two parts separated by a colon. The first part is a `cron`-style scheduling string. See the following web page for information about writing the string.

http://wiki.opensymphony.com/display/QRTZ1/CronTriggers+Tutorial

The second part of the `ibrix_at` command is the `ibrix_snap -A` command, which starts the snapshots. The following command takes a snapshot at 11:30am on the 10th of every month. The snapshot strategy, specified with the `-n` option, is `monthly1`.

```
ibrix_at "0 30 11 10 * ?" : ibrix_snap –A -f ifs1 -n monthly1
```

Other automated snapshot procedures

Use the following procedures to manage automated snapshots.

Modifying an automated snapshot scheme

A snapshot scheme can be modified only from the CLI. Use the following command:

```
ibrix_snap_strategy -e -n NAME -k KEEP -m MOUNT [-N NAMESPEC] [-M MOUNTSPEC]
```

Viewing automated snapshot schemes

On the GUI, you can view snapshot schemes on the Create Snapshot dialog box. Select **Recurring** as the Snapshot Type, and then select a snapshot scheme. A description of that scheme will be displayed.

To view all automated snapshot schemes or all schemes of a specific type using the CLI, execute the following command:

```
ibrix_snap_strategy -l [-T TYPE]
```

To see details about a specific automated snapshot scheme, use the following command:

```
ibrix_snap_strategy -i -n NAME
```

Deleting an automated snapshot scheme

A snapshot scheme can be deleted only from the CLI. Use the following command:

```
ibrix_snap_strategy -d -n NAME
```

Managing snapshots

This section describes how to manage snapshots using the CLI.

Creating a snapshot

Use the following command to create a snapshot:

```
<installdirectory>/bin/ibrix_snap -c -n SNAPFSNAME -f ORIGINFSNAME
```

For example, to create a snapshot named `ifs1_snap` for file system `ifs1`:

```
<installdirectory>/bin/ibrix_snap -c -n ifs1_snap -f ifs1
```

Mounting a snapshot

Include the `-M` option to the create command to automatically mount the snapshot file system after creating it. This makes the snapshot visible to authorized users. HP recommends that you do not allow writes to any snapshot file system.

```
<installdirectory>/bin/ibrix_snap -c –M –n SNAPFSNAME -f ORIGINFSNAME
```

For example, to create and mount a snapshot named `ifs1_snap` for file system `ifs1`:

```
<installdirectory>/bin/ibrix_snap -c -M –n ifs1_snap -f ifs1
```

Recovering system resources on snapshot failure

If a snapshot encounters insufficient resources when attempting to update its contents due to changes in the origin file system, the snapshot fails and is marked invalid. Data is no longer accessible in the snapshot. To clean up records in the configuration database for an invalid snapshot, use the following command:

```
<installdirectory>/bin/ibrix_snap -r -f SNAPFSLIST
```

For example, to clean up database records for a failed snapshot named `ifs1_snap`:

```
<installdirectory>/bin/ibrix_snap -r -f ifs1_snap
```

Deleting snapshots

Delete snapshots to free up resources when the snapshot is no longer needed or to create a new snapshot when you have already created the maximum allowed for your storage system.

```
<installdirectory>/bin/ibrix_snap -d -f SNAPFSLIST
```

For example, to delete snapshots `ifs0_snap` and `ifs1_snap`:

```
<installdirectory>/bin/ibrix_snap -d -f ifs0_snap,ifs1_snap
```

Viewing snapshot information

Use the following commands to view snapshot information from the CLI.

Listing snapshot information for all hosts

The `ibrix_snap -l` command displays snapshot information for all hosts. Sample output follows:

```
<installdirectory>/bin/ibrix_snap -l
```

```
NAME    NUM_SEGS  MOUNTED?  GEN  TYPE  CREATETIME
-----   --------  --------  ---  ----  ----------
snap1          3  No          6  msa   Wed Oct 7 15:09:50 EDT 2009
```

The following table lists the output fields for `ibrix_snap -l`.

Field	Description
NAME	Snapshot name.
NUM_SEGS	Number of segments in the snapshot.
MOUNTED?	Snapshot mount state.
GEN	Number of times the snapshot configuration has been changed in the configuration database.
TYPE	Snapshot type, based on the underlying storage system.
CREATETIME	Creation timestamp.

Listing detailed information about snapshots

Use the `ibrix_snap -i` command to monitor the status of active snapshots. You can use the command to ensure that the associated snapshot stores are not full.

```
<installdirectory>/bin/ibrix_snap -i
```

To list information about snapshots of specific file systems, use the following command:

```
<installdirectory>/bin/ibrix_snap -i [-f SNAPFSLIST]
```

The `ibrix_snap -i` command lists the same information as `ibrix_fs -i`, plus information fields specific to snapshots. Include the `-f SNAPFSLIST` argument to restrict the output to specific snapshot file systems.

The following example shows only the snapshot-specific fields. To view an example of the common fields, see "Viewing file system information" (page 26).

```
SEGMENT  OWNER     LV_NAME                 STATE            BLOCK_SIZE  CAPACITY(GB)
FREE(GB)  AVAIL(GB)  FILES  FFREE  USED%  BACKUP  TYPE   TIER  LAST_REPORTED
-------  --------  ---------------------  ---------------  ----------  ------------
--------  ---------  -----  -----  -----  ------  -----  ----  -------------
1        ib50-243  ilv11_msa_snap9__snap  OK, SnapUsed=4%       4,096          0.00
0.00       0.00      0      0      0            MIXED         7 Hrs 56 Mins 46 Secs ago
2        ib50-243  ilv12_msa_snap9__snap  OK, SnapUsed=6%       4,096          0.00
0.00       0.00      0      0      0            MIXED        7 Hrs 56 Mins 46 Secs ago
3        ib50-243  ilv13_msa_snap9__snap  OK, SnapUsed=6%       4,096          0.00
0.00       0.00      0      0      0            MIXED        7 Hrs 56 Mins 46 Secs ago
4        ib50-243  ilv14_msa_snap9__snap  OK, SnapUsed=8%       4,096          0.00
0.00       0.00      0      0      0            MIXED        7 Hrs 56 Mins 46 Secs ago
5        ib50-243  ilv15_msa_snap9__snap  OK, SnapUsed=6%       4,096          0.00
0.00       0.00      0      0      0            MIXED        7 Hrs 56 Mins 46 Secs ago
6        ib50-243  ilv16_msa_snap9__snap  OK, SnapUsed=5%       4,096          0.00
0.00       0.00      0      0      0            MIXED        7 Hrs 56 Mins 46 Secs ago
```

The following table lists the output fields for `ibrix_snap -i`.

Field	Description
SEGMENT	Snapshot segment number.
OWNER	The file serving node that owns the snapshot segment.
LV_NAME	Logical volume.
STATE	State of the snapshot.
BLOCK_SIZE	Block size used for the snapshot.
CAPACITY (GB)	Size of this snapshot file system, in GB.
FREE (GB)	Free space on this snapshot file system, in GB.
AVAIL (GB)	Space available for user files, in GB.
FILES	Number of files that can be created in this snapshot file system.
FFREE	Number of unused file inodes available in this snapshot file system.
USED%	Percentage of total storage occupied by user files.
BACKUP	Backup host name.
TYPE	Segment type. Mixed means the segment can contain both directories and files.
TIER	Tier to which the segment was assigned.
Last Reported	Last time the segment state was reported.

Accessing snapshot file systems

By default, snapshot file systems are mounted in two locations on the file serving nodes:

- */<snapshot_name>*
- */<original_file_system>/.<snapshot_name>*

For example, if you take a snapshot of the `fs1` file system and name the snapshot `fs1_snap1`, it will be mounted at `/fs1_snap1` and at `/fs1/.fs1_snap1`.

X9000 clients must mount the snapshot file system (*/<snapshot_name>*) to access the contents of the snapshot.

NFS and CIFS clients can access the contents of the snapshot through the original file system (such as `/fs1/.fs1_snap1`) or they can mount the snapshot file system (in this example, `/fs1_snap1`).

The following window shows an NFS client browsing the snapshot file system `.fs1_snap2` in the `fs1_nfs` file system.

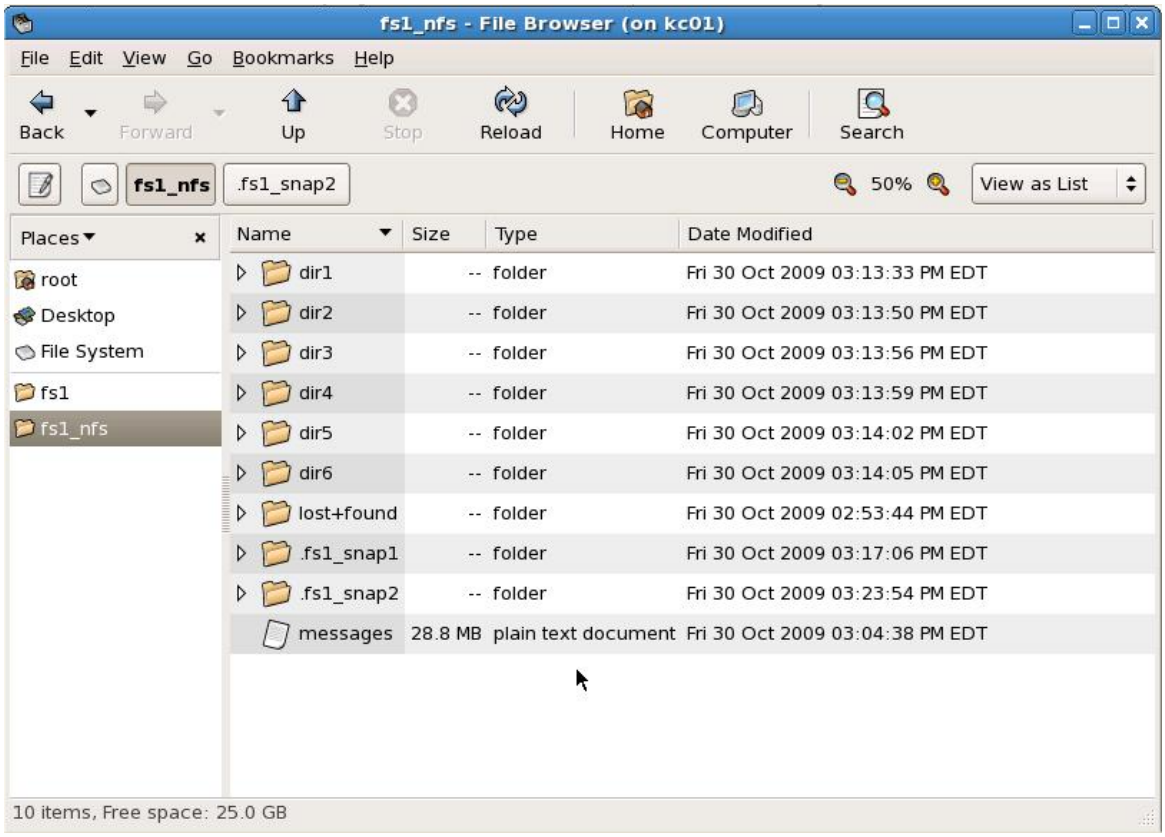
fs1_nfs - File Browser (on kc01)
File Edit View Go Bookmarks Help
Back Forward Up Stop Reload Home Computer Search
fs1_nfs .fs1_snap2 50% View as List
Places
root
Desktop
File System
fs1
fs1_nfs
Name Size Type Date Modified
dir1 -- folder Fri 30 Oct 2009 03:13:33 PM EDT
dir2 -- folder Fri 30 Oct 2009 03:13:50 PM EDT
dir3 -- folder Fri 30 Oct 2009 03:13:56 PM EDT
dir4 -- folder Fri 30 Oct 2009 03:13:59 PM EDT
dir5 -- folder Fri 30 Oct 2009 03:14:02 PM EDT
dir6 -- folder Fri 30 Oct 2009 03:14:05 PM EDT
lost+found -- folder Fri 30 Oct 2009 02:53:44 PM EDT
.fs1_snap1 -- folder Fri 30 Oct 2009 03:17:06 PM EDT
.fs1_snap2 -- folder Fri 30 Oct 2009 03:23:54 PM EDT
messages 28.8 MB plain text document Fri 30 Oct 2009 03:04:38 PM EDT
10 items, Free space: 25.0 GB

The next window shows a CIFS client accessing the snapshot file system `.fs1_snap1`. The original file system is mapped to drive X.

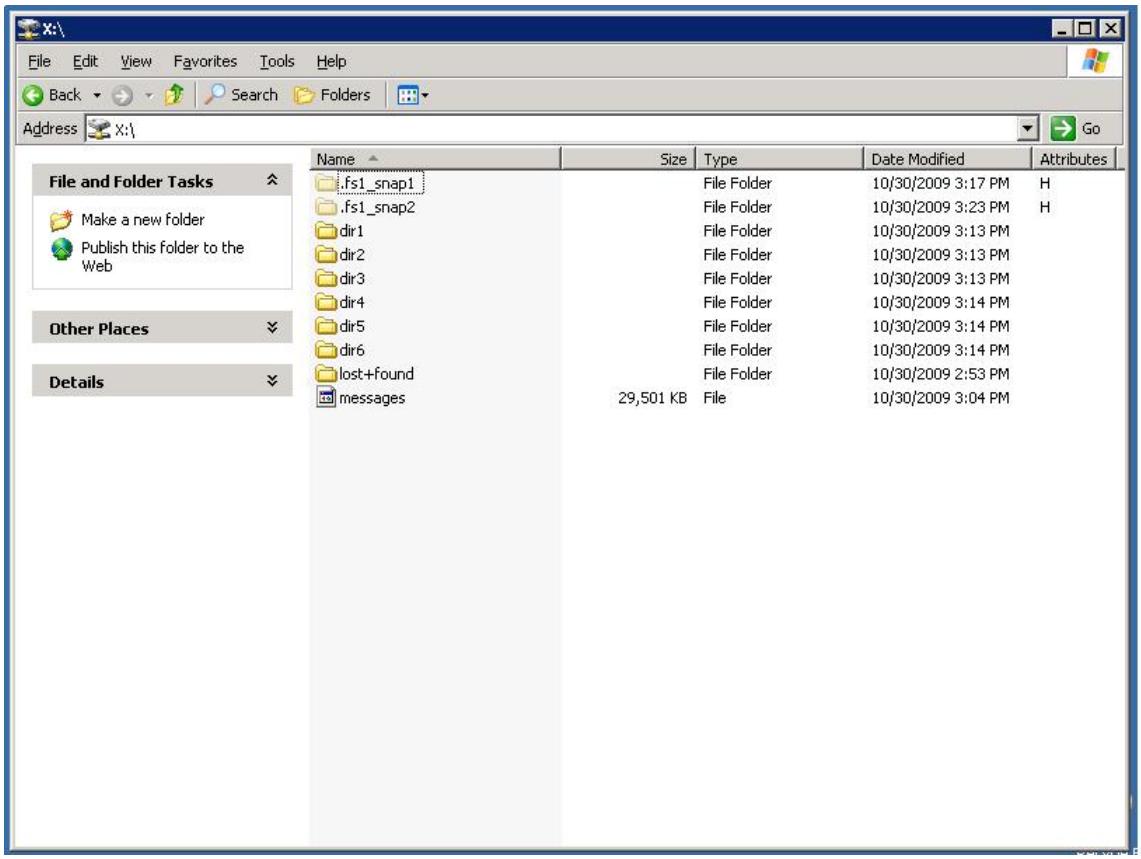
X:\
File Edit View Favorites Tools Help
Back Search Folders
Address X:\ Go
File and Folder Tasks
Make a new folder
Publish this folder to the Web
Other Places
Details
Name Size Type Date Modified Attributes
.fs1_snap1 File Folder 10/30/2009 3:17 PM H
.fs1_snap2 File Folder 10/30/2009 3:23 PM H
dir1 File Folder 10/30/2009 3:13 PM
dir2 File Folder 10/30/2009 3:13 PM
dir3 File Folder 10/30/2009 3:13 PM
dir4 File Folder 10/30/2009 3:14 PM
dir5 File Folder 10/30/2009 3:14 PM
dir6 File Folder 10/30/2009 3:14 PM
lost+found File Folder 10/30/2009 2:53 PM
messages 29,501 KB File 10/30/2009 3:04 PM

Troubleshooting snapshots

Snapshot reserve is full and the MSA2000 is deleting snapshot volumes

When the snapshot reserve is full, the MSA2000 will delete snapshot volumes on the storage array, leaving the device entries on the file serving nodes. To correct this situation, take the following steps:

1. Stop I/O or any applications that are reading or writing to the snapshot file systems.
2. Log on to the active management console.
3. Unmount all snapshot file systems.
4. Delete all snapshot file systems to recover space in the snapshot reserve.

CIFS clients receive an error when creating a snapshot

CIFS clients might see the following error when attempting to create a snapshot:

```
Make sure you are connected to the network and try again
```

This error is generated when the snapshot creation takes longer than the CIFS timeout, causing the CIFS client to determine that the server has failed or the network is disconnected. To avoid this situation, do not take snapshots during periods of high CIFS activity.

13 Using data tiering

This chapter describes how to configure and manage data tiering.

Overview

Use the data tiering feature to set a preferred tier where newly created files will be stored. Once files are created, you can use a tiering job to move them from initial storage, based on file attributes. When you start the tiering job, you can specify a desired number of file replicas to add (default is 0).

- You can use any naming convention you choose to identify each tier grouping.
- You can create as many named tiers as your environment requires.
- A segment can belong to only one tier at a time, but it can be reassigned at any time.

Assigning file system segments to distinct tiers allows you to categorize areas in your file system and storage network. After the tiers are created, you can use rule-based commands to automatically move files from tier to tier, based on file attributes.

For example, a storage network composed of newer and older drives might be tiered based on the access speed of the storage media. Using descriptive tier names such as "fast" and "slow," you assign each segment in your file system to the appropriate tier, creating a tier of fast segments and a tier of slower segments.

Using the migrator utility, you can write and execute rules that move frequently accessed files into the fast storage tier, and move files that are rarely accessed to less expensive storage media in the slow tier. This is just one example of moving files based on a tiering structure. Other examples might be based on the type of file being stored, such as storing all streaming files in a tier or moving all files over a certain size to a specific tier.

IMPORTANT: Data tiering has a cool-down period of approximately 10 minutes. If a file was last accessed during the cool-down period, the file will not be moved.

Moving files between tiers

X9000 Software allows you to set a preferred tier where newly created files will be stored. Once files are created, you can use a tiering job to move them from initial storage, based on file attributes.

Tiering is the ability to programmatically move files from one tier to another in the same file system. You write rules based on file attributes such as modification time, access time, file size, or file type to define the tiering policy. The policy then determines which files are to be moved and when.

- A tiering policy (a set of rules) applies to individual files in a specific file system.
- If a file meets the criteria of a rule, it will be moved from its current tier to the rule's target tier.
- A tiering policy (once configured) is executed via a command or as a `cron` job, and is performed in the background by the system.
- Files created or changed within the last five minutes are considered active and will not be moved.

Writing a rule to implement a policy

A *tiering policy* is defined by rules. Each rule identifies file attributes to match. It also specifies the source tier to scan and the destination tier where files that meet the rule's criteria will be moved and stored.

Note the following:

- A tiering policy can contain multiple rules.
- Tiering rules are based on individual file attributes.
- All rules are executed when the tiering policy is applied during execution of the `ibrix_migrator` command.
- It is important that different rules do not target the same files, especially if different destination tiers are specified. If tiering rules are ambiguous, the final destination for a file is not predictable. See "Ambiguous rules" (page 114), for more information.

The following are examples of attributes that can be specified in rules. All attributes are listed in "Rule keywords" (page 113). You can use AND and OR operators to create combinations.

Access time

- File was last accessed *x* or more days ago
- File was accessed in the last *y* days

Modification time

- File was last modified *x* or more days ago

File size—greater than *n* K

File name or File type—`jpg`, `wav`, `exe` (include or exclude)

File ownership—owned by user(s) (include or exclude)

Use of the tiering assignments or policy on any file system is optional. Tiering is not assigned by default; there is no "default" tier.

Tiered file systems

You can create a tiered file system at any time. You can assign tier names to file serving nodes when creating a file system or assign tiers after the file system is in place.

- Use the `ibrix_fs -t` *`TIERNAME`* option to create or expand a file system with tiered segments.
- Use `ibrix_tier` to assign segments to tiers, change a segment's tier assignment, and delete a tier.
- Use `ibrix_migrator` to start and stop tiering operations, and to write and manage the rules that govern tiering policy.

You can configure and manage data tiering from the management console GUI or from the CLI using the `ibrix_tier` and `ibrix_migrator` commands. This section provides instructions for using the CLI commands. For more information about these commands, see the *HP X9000 File Serving Software CLI Reference Guide*.

Creating a file system that uses data tiering

The `ibrix_fs` command, which is used to create or expand a file system, includes a `-t` *`TIERNAME`* option that can be used to create a data tier. Use the following command to create the file system:

`<installdirectory>``/bin/ibrix_fs -c -f` *`FSNAME`* `-s` *`LVLIST`* `-t` *`TIERNAME`* `....`

The command creates file system *`FSNAME`* with the segments in *`LVLIST`* assigned to *`TIERNAME`*. *`TIERNAME`* can be any alphanumeric, case-sensitive, text string. Tier assignment is not affected by other options that can be set with the `ibrix_fs` command.

NOTE: A tier is created whenever a segment is assigned to it. Be careful to spell the name of the tier correctly when you add segments to an existing tier. If you make an error in the name, a new tier is created with the incorrect tier name, and no error is recognized.

Expanding a file system

You can make tier assignments for new segments as you add them.

NOTE: Verify that no tiering job is running before executing the file system expansion commands. The expansion takes priority and the tiering job is terminated.

To including tiering at the time you expand a file system, use the following command:

```
<installdirectory>/bin/ibrix_fs -e -f FSNAME -s LVLIST -t TIERNAME
```

This command extends the file system *FSNAME* with the segments in *LVLIST* and assigns the segments to *TIERNAME*. If tiering policy rules are already defined for this file system, the `-t` option is required.

To extend file system *FSNAME* with the listed tiered segment/owner pairs:

```
<installdirectory>/bin/ibrix_fs -e -f FSNAME -S LV1:HOSTNAME1,LV2:HOSTNAME2,...
-t TIERNAME
```

To extend file system *FSNAME* with the physical volumes in *PVLIST*:

```
<installdirectory>/bin/ibrix_fs -e -f FSNAME -p PVLIST -t TIERNAME
```

Allocation policy

The allocation policy manager is responsible for determining the initial placement of files when they are created. The allocation policy command can use all of the segments currently in the tier.

To specify that all segments within a named tier be used for initial placement of new files, use the `ibrix_fs_tune` command to explicitly prefer the tier.

```
<installdirectory>/bin/ibrix_fs_tune -f FSNAME {-h HOSTLIST|-g GROUPLIST} -t TIERNAME
```

This command designates a primary tier and prefers the pool of segments identified by the tier name (*TIERNAME*) on a list of specified hosts (*HOSTLIST*) or hostgroups (*GROUPLIST*). You can prefer only one primary tier in a file system.

Unpreferred segments in other secondary tiers are ignored by the allocation policy manager and only used by `ibrix_migrator` when moving files according to the tiering policy. For example, in a file system with six segments, you might specify the following configuration:

1. Add the first three segments to tier T1.

   ```
   <installdirectory>/bin/ibrix_tier -a -f ifs1 -t T1 -S 1,2,3
   ```

2. Execute `ibrix_fs_tune` to prefer tier T1, which contains segments `s1-s3` for initial file writes.

   ```
   <installdirectory>/bin/ibrix_fs_tune -f ifs1 -h ss2 -t T1
   ```

3. Assign segments `s4-s6` to a second tier.

   ```
   <installdirectory>/bin/ibrix_tier -a -f ifs1 -t T2 -S 4,5,6
   ```

The tiering policy can now be used to move matching files from tier T1, where they will be stored on creation, to tier T2 based on file attribute conditions.

If you add segments to the primary tier at a later time, the segments automatically become preferred. Similarly, the data tiering feature unprefers any segment removed from the primary tier at a later time.

Managing a tiered file system and tiering policy

After tiers have been defined on a file system, you can:

- Add segments to existing tiers.
- Add or delete tiers.
- Add or delete rules for your tiering policy.
- Run and monitor your tiering policy.

You can perform these management activities from either the management console GUI or the CLI. This section provides instructions for managing via the CLI.

Assigning segments to tiers

Segments can be assigned to tiers when a file system is created or expanded, or at any time when a tiering policy job is not running. Similarly, tier assignments can be changed or removed at any time using the `ibrix_tier` command, provided there are no tiering jobs running.

To manually assign one or more segments to a tier, use the following command:

```
<installdirectory>/bin/ibrix_tier -a -f FSNAME -t TIERNAME -S SEGLIST
```

NOTE: A tier is created whenever a segment is assigned to it. Be careful to spell the name of the tier correctly when you add segments to an existing tier. If you make an error in the name, a new tier is created with the incorrect tier name, and no error is recognized.

Removing tier assignments

To remove one or more segments from their assigned tier, use the following command:

```
<installdirectory>/bin/ibrix_tier -u -f FSNAME [-S SEGLIST]
```

If no segment list is supplied, all segments in the named file system are unassigned. A tier is automatically deleted when all segments are removed from it. If there are no segments, there is no tier.

Deleting a tier

Use the `ibrix_tier` command to delete a tier. This step automatically unassigns all segments in that tier without reassigning them. Before deleting a tier, you must:

- Delete all policy rules defined for the tier.
- Allow any active tiering jobs to complete.

To delete a tier, use the following command:

```
<installdirectory>/bin/ibrix_tier -d -f FSNAME -t TIERNAME
```

Listing tier information

To view information about tiers on a specific file system or on all file systems, use the following command:

```
<installdirectory>/bin/ibrix_tier -l -f FSNAME [-t TIERNAME]
```

The `-l` option lists the tiers associated with *FSNAME*. Specify *TIERNAME* to view information about a specific tier.

Listing tiering policy information

All of the tiering rules defined on a file system form the file system's tiering policy. Use the following command to view the policy. To view rule definitions on a specific file system or on all file systems, include the `-r` option.

```
<installdirectory>/bin/ibrix_migrator -v -r [-f FSNAME]
```

For example:

```
<installdirectory>/bin/ibrix_migrator -v -r -f ifs2
```

The output lists the file system name, the rule ID (IDs are assigned in the order in which rules are added to the configuration database), the rule definition, and the source and destination tiers. For example, if you specify the following rule:

```
ibrix_migrator -A -f ifs2 -r 'mtime older than 1 month and
( name = "*.jpg" or name = "*.gif" )' -S T1 -D T2
```

The output lists the following:

```
ifs2 2 mtime older than 1 month and ( name = "*.jpg" or name = "*.gif" ) T1 T2
```

The second column in the output lists the ID number assigned to the rule. In the example above, the ID number is 2.

Deleting a tiering policy rule

Before deleting a rule, view the policy information, as described in "Listing tiering policy information" (page 111), and obtain the rule's ID number. The ID number is required in the delete command.

To delete a rule, use the following command:

```
<installdirectory>/bin/ibrix_migrator -d -f FSNAME -r RULE_ID
```

RULE_ID is the rule's ID number (2 in the example).

```
<installdirectory>/bin/ibrix_migrator -d -f ifs2 -r 2
```

Starting and stopping a tiering operation

Once a tiering policy is defined, tiering operations can be started and stopped using the `ibrix_migrator` command. Only one tiering operation can run on a file system at any time. Tiering operations are treated as run-to-completion tasks that are not restarted on failure, and cannot be paused and later resumed. However, tiering can be started if a server is in the InFailover state.

NOTE: The `ibrix_migrator` command cannot be run at the same time as `ibrix_rebalance`.

To start a tiering operation, use the following command:

```
<installdirectory>/bin/ibrix_migrator -s -f FSNAME
```

To stop a tiering operation, use the following command:

```
<installdirectory>/bin/ibrix_migrator -k -t TASKID [-F]
```

Reviewing tiering job status

To view tiering tasks, including the task state, use the following command:

```
<installdirectory>/bin/ibrix_migrator -i [-f FSNAME]
```

Writing tiering rules

A tiering policy consists of one or more rules, each identifying a desired movement of files between tiers. You can write rules using the management console GUI, or you can write them directly to the configuration database using the `ibrix_migrator -A` command.

This section provides definitions of rule components and examples of rules.

Operators and date/time qualifiers

Valid rules operators are `<`, `<=`, `=`, `!=`, `>`, `>=`, and boolean `and` and `or`.

Use the following qualifiers for fixed times and dates:

- **Time:** Enter as three pairs of colon-separated integers using a 24-hour clock. The format is *hh:mm:ss* (for example, **15:30:00**).
- **Date:** Enter as *yyyy-mm-dd [hh:mm:ss]*, where time of day is optional (for example, **2008-06-04** or **2008-06-04 15:30:00**). Note the space separating the date and time.

When specifying an absolute date and/or time, the rule must use a compare type operator (< | <= | = | != | > | >=). For example:

```
ibrix_migrator -A -f ifs2 -r "atime > '2010-09-23' " -S TIER1 -D TIER2
```

Use the following qualifiers for relative times and dates:

- **Relative time:** Enter in rules as `year` or `years`, `month` or `months`,`week` or `weeks`, `day` or `days`, `hour` or `hours`.
- **Relative date:** Use `older than` or `younger than`. The rules engine uses the time the `ibrix_migrator` command starts execution as the start time for the rule. It then computes the required time for the rule based on this start time. For example, **`ctime older than 4 weeks`** refers to that time period more that 4 weeks before the start time.

The following example uses a relative date:

```
ibrix_migrator -A -f ifs2 -r "atime older than 2 days " -S TIER1 -D
TIER2
```

Rule keywords

The following keywords can be used in rules.

Keyword	Description
`atime`	Access time, used in a rule as a fixed or relative time.
`ctime`	Change time, used in a rule as a fixed or relative time.
`mtime`	Modification time, used in a rule as a fixed or relative time
`gid`	An integer corresponding to a group ID.
`gname`	A string corresponding to a group name. Enclose the name string in double quotes.
`uid`	An integer corresponding to a user ID.
`uname`	A string corresponding to a user name, where the user is the owner of the file. Enclose the name string in double quotes.
`type`	File system entity the rule operates on. Currently, only the file entity is supported.
`size`	In size-based rules, the threshold value for determining migration. Value is an integer specified in K (KB), M (MB), G (GB), and T (TB). Do not separate the value from its unit (for example, 24K).
`name`	Regular expression. A typical use of a regular expression is to match file names. Enclose a regular expression in double quotes. The * wildcard is valid (for example, `name = "*.mpg"`). A name cannot contain a / character. You cannot specify a path; only a filename is allowed.
`path`	Path name that allows these wild cards: `*`, `?`, `/`. For example, if the mountpoint for the file system is `/mnt`, `path=ibfs1/mydir/*` matches the entire directory subtree under `/mnt/ibfs1/mydir`. (A path cannot start with a `/`).
`strict_path`	Path name that rigidly conforms to UNIX shell file name expansion behavior. For example, `strict_path=/mnt/ibfs1/mydir/*` matches only the files that are explicitly in the `mydir` directory, but does not match any files in subdirectories of `mydir`.

Migration rule examples

When you write a rule, identify the following components:

- File system (`-f`)
- Source tier (`-S`)
- Destination tier (`-D`)

Use the following command to write a rule. The rule portion of the command must be enclosed in single quotes.

```
ibrix_migrator -A -f FSNAME -r 'RULE' -S SOURCE_TIER -D DEST_TIER
```

Examples:

The rule in the following example is based on the file's last modification time, using a relative time period. All files whose last modification date is more than one month in the past are moved.

```
# ibrix_migrator -A -f ifs2 -r 'mtime older than 1 month' -S T1 -D T2
```

In the next example, the rule is modified to limit the files being migrated to two types of graphic files. The `or` expression is enclosed in parentheses, and the `*` wildcard is used to match filename patterns.

```
# ibrix_migrator -A -f ifs2 -r 'mtime older than 1 month and
( name = "*.jpg" or name = "*.gif" )' -S T1 -D T2
```

In the next example, three conditions are imposed on the migration. Note that there is no space between the integer and unit that define the size threshold (`10M`):

```
# ibrix_migrator -A -f ifs2 -r 'ctime older than 1 month and type = file
and size >= 10M' -S T1 -D T2
```

The following example uses the `path` keyword. It moves files greater than or equal to 5M that are under the directory `/ifs2/tiering_test` from TIER1 to TIER2:

```
ibrix_migrator -A -f ifs2 -r "path = tiering_test and size >= 5M" -S
TIER1 -D TIER2
```

Rules can be group- or user-based as well as time- or data-based. In the following example, files associated with two users are migrated to `T2` with no consideration of time. The names are quoted strings.

```
# ibrix_migrator -A -f ifs2 -r 'type = file and ( uname = "ibrixuser"
or uname = "nobody" )' -S T1 -D T2
```

Conditions can be combined with `and` and `or` to create very precise tiering rules, as shown in the following example.

```
# ibrix_migrator -A -f ifs2 -r ' (ctime older than 3 weeks and ctime younger
than 4 weeks) and type = file and ( name = "*.jpg" or name = "*.gif" )
and (size >= 10M and size <= 25M)' -S T1 -D T2
```

Ambiguous rules

It is possible to write a set of ambiguous rules, where different rules could be used to move a file to conflicting destinations. For example, if a file can be matched by two separate rules, there is no guarantee which rule will be applied in a tiering job.

Ambiguous rules can cause a file to be moved to a specific tier and then potentially moved back. Examples of two such situations follow.

Example 1:

In the following example, if a `.jpg` file older than one month exists in tier 1, then the first rule moves it to tier 2. However, once it is in tier 2, it is matched by the second rule, which then moves the file back to tier 1.

```
# ibrix_migrator -A -f ifs2 -r ' mtime older than 1 month ' -S T1 -D T2
# ibrix_migrator -A -f ifs2 -r ' name = "*.jpg" ' -S T2 -D T1
```

There is no guarantee as to the order in which the two rules will be executed; therefore, the final destination is ambiguous because multiple rules can apply to the same file.

Example 2:

Rules can cause data movement in both directions, which can lead to issues. In the following example, the rules specify that all `.doc` files in tier 1 to be moved to tier 2 and all `.jpg` files in tier 2 be moved to tier 1. However, this might not succeed, depending on how full the tiers are.

```
# ibrix_migrator -A -f ifs2 -r ' name = "*.doc" ' -S T1 -D T2
# ibrix_migrator -A -f ifs2 -r ' name = "*.jpg" ' -S T2 -D T1
```

For example, if tier 1 is filled with `.doc` files to 70% capacity and tier2 is filled with `.jpg` files to 80% capacity, then tiering might terminate before it is able to fully "swap" the contents of tier 1 and tier 2. The files are processed in no particular order; therefore, it is possible that more `.doc` files will be encountered at the beginning of the job, causing space on tier 2 to be consumed faster than on tier 1. Once a destination tier is full, obviously no further movement in that direction is possible.

These rules in these two examples are ambiguous because they give rise to possible conflicting file movement. It is the user's responsibility to write unambiguous rules for the data tiering policy for their file systems.

14 Using file allocation

This chapter describes how to configure and manage file allocation.

Overview

X9000 Software allocates new files and directories to segments according to the allocation policy and segment preferences that are in effect for a client. An allocation policy is an algorithm that determines the segments that are selected when clients write to a file system.

File allocation policies

Different allocation policies can be set for files and directories. The following table lists the available allocation policies.

By default, file systems use the RANDOM allocation policy, and all segments are preferred. These defaults are adequate for most cluster environments. In some cases, HP Support may recommend that you change the defaults as follows to optimize file storage for your system:

- Change the default file and directory allocation policy.
- Prefer a pool of segments for storing all new files and directories.
- Declare the first segment where a file or directory will be stored.

IMPORTANT: Changing the default allocation policy and segment preferences will alter file system storage behavior. Contact HP Support before changing either of these defaults.

Although the allocation policy and segment preferences are set locally on X9000 clients, they are set on the NFS or CIFS server for NFS or CIFS clients. X9000 clients write directly to file serving nodes according to the file allocation settings in effect for them. In contrast, NFS and CIFS clients write to a file system only via their NFS or CIFS server; therefore, the allocation policy and segment preferences that are in effect on the NFS or CIFS server determine the segments where the data is stored.

To choose a segment, an X9000 client or NFS/CIFS server evaluates the allocation policy and segment preferences that are defined for it.

The following table lists the file allocation policies.

Policy	Description
AUTOMATIC	Lets the X9000 Software select the allocation policy.
DIRECTORY	Allocates files to the segment where its parent directory is located.
LOCAL	Allocates files to a file serving node's local segments.
RANDOM	Allocates files to a randomly chosen segment among preferred segments. This is the default policy.
ROUNDROBIN	Allocates files to preferred segments in segment order, returning to the first segment (or the designated starting segment) when a file has been allocated to the last segment.
STICKY	Allocates files to one segment until the segment's storage limit is reached, and then moves to the next segment as determined by the AUTOMATIC file allocation policy.
NONE	Sets the directory allocation policy only. Causes the directory allocation policy to revert to its default, which is the policy set for file allocation. Use NONE only to set file and directory allocation to the same policy.
HOST_ROUNDROBIN_NB	For clusters with more than 16 file serving nodes, takes a subset of the servers to be used for file creation and rotates this subset on a regular periodic basis. This policy should be used only under the direction of HP Support.

How file allocation settings are evaluated

X9000 clients and NFS/CIFS clients use the following precedence rules to evaluate the file allocation settings that are in effect for them:

- **If the host uses the default file and directory allocation policies and segment preferences:** The RANDOM policy is applied, and a segment is chosen from among all segments.
- **If the host uses a nondefault file and/or directory allocation policy and the default segment preferences:** Only the file or the directory allocation policy is applied, and a segment is chosen from among the available segments.
- **If the host uses a nondefault file and/or directory allocation policy and nondefault segment preferences:** A segment is chosen according to the following rules:
  a. From the pool of preferred segments, select a segment according to the allocation policy set for the client, and store the file in that segment if there is room. If all segments in the pool are full, proceed to the next rule.
  b. Use the AUTOMATIC allocation policy to choose a segment with enough storage room from among the available segments, and store the file there.

Using hostgroups for file allocation settings

The file allocation commands include an option, `-g GROUPLIST`, that specifies one or more hostgroups. This is a convenient way to configure file allocation settings for a set of X9000 clients. To configure file allocation settings for all X9000 clients, specify the `clients` hostgroup.

When file allocation settings take effect on X9000 clients

Although file allocation settings are executed immediately on file serving nodes, for X9000 clients, a file allocation intention is stored on the management console. When X9000 Software services start on a client, the client queries the management console for the file allocation settings that it should use and then implements them. If the services are already running on a client, you can force the client to query the management console by executing `ibrix_client` or `ibrix_lwhost --a` on the client, or by rebooting the client.

Guidelines for using file allocation CLI commands

Follow these guidelines when using CLI commands to perform any file allocation configuration tasks:

- To perform a task for NFS/CIFS clients, specify NFS/CIFS servers for the `-h HOSTLIST` argument.
- To perform a task for X9000 clients, either specify individual clients for the `-h HOSTLIST` argument or specify a hostgroup for the `-g GROUPLIST` argument.

Setting file and directory allocation policies

You can set a nondefault file or directory allocation policy for NFS/CIFS servers and X9000 clients. Before using the CLI commands to do this, see "Guidelines for using file allocation CLI commands" (page 117).

You can specify the first segment where the policy should be applied, but in practice this is useful only for the ROUNDROBIN policy.

Allocation policy names are case sensitive and must be entered as uppercase letters (for example, RANDOM).

NOTE: If your file and directory allocation policies are different and you want to make them the same, first verify that the file allocation policy is set correctly. Next, set the directory allocation policy to NONE. The directory allocation policy then reverts to its default value, which is the policy set for file allocation.

Setting a file allocation policy

To set a file allocation policy, use the following command:

```
<installdirectory>/bin/ibrix_fs_tune -f FSNAME {-h HOSTLIST|-g GROUPLIST}
-s LVNAMELIST -p POLICY [-S STARTSEGNUM]
```

For example, to set the ROUNDROBIN policy for files only on NFS-exported file system `ifs1` on file serving node s1.hp.com, starting at segment `ilv1`:

```
<installdirectory>/bin/ibrix_fs_tune -f ifs1 -h s1.hp.com -p ROUNDROBIN -s ilv1
```

Setting a directory allocation policy

To set a directory allocation policy, use the following command. The `-R` argument declares that the policy is for directories only.

```
<installdirectory>/bin/ibrix_fs_tune -f FSNAME {-h HOSTLIST|-g GROUPLIST}
-p POLICY [-S STARTSEGNUM] [-R]
```

For example, to set the ROUNDROBIN directory allocation policy on NFS-exported file system `ifs1` for file serving node s1.hp.com, starting at segment `ilv1`:

```
<installdirectory>/bin/ibrix_fs_tune -f ifs1 -h s1.hp.com -p ROUNDROBIN -R
```

Setting segment preferences

There are two ways to prefer segments for X9000 clients, NFS/CIFS servers, or hostgroups:

- Prefer a pool of segments for the clients to use.
- Prefer a single segment for files created by a specific user or group on the clients.

Both methods can be in effect at the same time. For example, you can prefer a segment for a user and then prefer a pool of segments for the clients on which the user will be working.

If you have set segment preferences and need to change them later on, use the `ibrix_fs_tune` command to specify the new segment preferences.

Creating a pool of preferred segments

A segment pool can consist of individually selected segments, all segments local to an NFS/CIFS server, or all segments. Clients will apply the allocation policy that is in effect for them to choose a segment from the segment pool.

NOTE: Segments are always created in the preferred condition. If you want to have some segments preferred and others unpreferred, first select a single segment and prefer it. This action *unprefers* all other segments. You can then work with the segments one at a time, preferring and unpreferring as required. By design, the system cannot have zero preferred segments. If only one segment is preferred and you unprefer it, *all* segments become preferred.

When preferring multiple pools of segments (for example, one for X9000 clients and one for NFS/CIFS clients), make sure that no segment appears in both pools.

Use the following command to specify the pool by logical volume name (*LVNAMELIST*):

```
<installdirectory>/bin/ibrix_fs_tune -f FSNAME {-h HOSTLIST|-g GROUPLIST} -s LVNAMELIST
```

Use the following command and the `LOCAL` keyword to create a pool of all segments on NFS/CIFS servers. Use the `ALL` keyword to restore the default segment preferences (for more information see "Restoring the default segment preference" (page 119)).

```
<installdirectory>/bin/ibrix_fs_tune -f FSNAME {-h HOSTLIST|-g GROUPLIST}
-S {SEGNUMLIST|ALL|LOCAL}
```

Restoring the default segment preference

The default is for all file system segments to be preferred. Use the following command to restore the default value:

```
<installdirectory>/bin/ibrix_fs_tune -f FSNAME {-h HOSTLIST|-g GROUPLIST} -S ALL
```

Tuning allocation policy settings

To optimize system performance, you can globally change the following allocation policy settings for a file system:

- File allocation policy (see "File allocation policies" (page 116) for a list of available policies).
- Starting segment number for applying changes.
- Preallocation: Number of KB to preallocate for files. The default is a preallocated file size of 256 KB.
- Readahead: Number of KB in a file to pre-fetch. The default is 128 KB.
- NFS readahead: Number of KB in a file to pre-fetch on NFS systems. The default is 128 KB.

NOTE: Before tuning allocation policy settings, contact HP Support for guidance on selecting the best values for your installation.

Use the following command to restore the default file allocation policy:

```
<installdirectory>/bin/ibrix_fs_tune -f FSNAME {-h HOSTLIST|-g GROUPLIST} -p -U
```

Listing allocation policies

Use the following command to list the preferred segments (the `-S` option) or the allocation policy (the `-P` option) for the specified hosts, hostgroups, or file system.

```
<installdirectory>/bin/ibrix_fs_tune -l [-S] [-P] [-h HOSTLIST | –g GROUPLIST] [-f FSNAME]

HOSTNAME          FSNAME    POLICY    STARTSEG   DIRPOLICY   DIRSEG
   SEGBITS   READAHEAD    PREALLOC  HWM        SWM
mak01.hp.com      ifs1      RANDOM    0           NONE         0
   DEFAULT   DEFAULT      DEFAULT   DEFAULT    DEFAULT
```

15 Support and other resources

Contacting HP

For worldwide technical support information, see the HP support website:

http://www.hp.com/support

Before contacting HP, collect the following information:

- Product model names and numbers
- Technical support registration number (if applicable)
- Product serial numbers
- Error messages
- Operating system type and revision level
- Detailed questions

Related information

The following documents provide related information:

- *HP StorageWorks X9000 File Serving Software Release Notes*
- *HP StorageWorks X9000 File Serving Software CLI Reference*
- *HP StorageWorks X9300 Network Storage Gateway Administrator Guide*
- *HP StorageWorks X9320 Network Storage System Administrator Guide*
- *HP StorageWorks X9720 Network Storage System Administrator Guide*
- *HP StorageWorks X9000 File Serving Software Installation Guide*

Related documents are available on the Manuals page at http://www.hp.com/support/manuals. On the Manuals page, select **storage** > **NAS Systems** > **NAS/Storage Servers** > **HP StorageWorks X9000 Network Storage Systems**.

HP websites

For additional information, see the following HP websites:

- http://www.hp.com
- www.hp.com/go/x9000
- http://www.hp.com/go/storage
- http://www.hp.com/support/manuals

Subscription service

HP recommends that you register your product at the Subscriber's Choice for Business website:

http://www.hp.com/go/e-updates

After registering, you will receive e-mail notification of product enhancements, new driver versions, firmware updates, and other product resources.

Glossary

ACE	access control entry.
ACL	access control list.
ADS	Active Directory Service.
ALB	Advanced load balancing.
BMC	Baseboard Management Configuration.
CIFS	Common Internet File System. The protocol used in Windows environments for shared folders.
CLI	Command-line interface. An interface comprised of various commands which are used to control operating system responses.
CSR	Customer self repair.
DAS	Direct attach storage. A dedicated storage device that connects directly to one or more servers.
DNS	Domain name system.
FTP	File Transfer Protocol.
GSI	Global service indicator.
HA	High availability.
HBA	Host bus adapter.
HCA	Host channel adapter.
HDD	Hard disk drive.
IAD	HP X9000 Software Administrative Daemon.
iLO	Integrated Lights-Out.
IML	Initial microcode load.
IOPS	I/Os per second.
IPMI	Intelligent Platform Management Interface.
JBOD	Just a bunch of disks.
KVM	Keyboard, video, and mouse.
LUN	Logical unit number. A LUN results from mapping a logical unit number, port ID, and LDEV ID to a RAID group. The size of the LUN is determined by the emulation mode of the LDEV and the number of LDEVs associated with the LUN.
MTU	Maximum Transmission Unit.
NAS	Network attached storage.
NFS	Network file system. The protocol used in most UNIX environments to share folders or mounts.
NIC	Network interface card. A device that handles communication between a device and other devices on a network.
NTP	Network Time Protocol. A protocol that enables the storage system's time and date to be obtained from a network-attached server, keeping multiple hosts and storage devices synchronized.
OA	HP Onboard Administrator.
OFED	OpenFabrics Enterprise Distribution.
OSD	On-screen display.
OU	Active Directory Organizational Units.
RO	Read-only access.
RPC	Remote Procedure Call.
RW	Read-write access.
SAN	Storage area network. A network of storage devices available to one or more servers.
SAS	Serial Attached SCSI.

SELinux	Security-Enhanced Linux.
SFU	Microsoft Services for UNIX.
SID	Secondary controller identifier number.
SNMP	Simple Network Management Protocol.
TCP/IP	Transmission Control Protocol/Internet Protocol.
UDP	User Datagram Protocol.
UFM	Voltaire's Unified Fabric Manager client software.
UID	Unit identification.
USM	SNMP User Security Model.
VACM	SNMP View Access Control Model.
VC	HP Virtual Connect.
VIF	Virtual interface.
WINS	Windows Internet Naming Service.
WWN	World Wide Name. A unique identifier assigned to a Fibre Channel device.
WWNN	World wide node name. A globally unique 64-bit identifier assigned to each Fibre Channel node process.
WWPN	World wide port name. A unique 64-bit address used in a FC storage network to identify each device in a FC network.

Index

Symbols

/etc/likewise/vhostmap file, 63
32-bit mode, disable, 32
64-bit mode, enable, 11

A

authentication methods
  Active Directory, 42
  Local Users, 42
automated snapshots
  create from CLI, 101
  create on GUI, 98
  delete snapshot scheme, 103
  modify snapshot scheme, 103
  schedule from CLI, 102
  snapshot scheme, 97
  view snapshot scheme from CLI, 103

C

CIFS
  activity statistics per node, 49
  authentication methods, 42
  configure local groups, 45
  configure local users, 47
  configure nodes, 49
  file shares, 52
  Linux static user mapping, 59
  locking behavior, 64
  permissions management, 64
  RFC2037 support, 59
  SMB server consolidation, 62
  SMB signing, 51
  start or stop CIFS service, 49
CIFS shares
  add a share, 52
  add with MMC, 57
  delete a share, 54
  delete with MMC, 59
  manage with MMC, 54
  modify a share, 54
  view share information, 50
clients
  SAN I/O, 9
cluster
  components, 8
contacting HP, 120

D

data tiering
  allocation policy, 110
  assign segments to tiers, 111
  create tiered file system, 109
  delete a tier, 111
  delete tier policy rule, 112
  expand tiered file system, 110
  list tasks, 112
  list tier information, 111
  list tier policy information, 111
  move files, 108
  remove segments from tiers, 111
  start or stop operation, 112
  tiering policy rules, 108
  write tiering rules, 112
document
  related documentation, 120

E

Export Control, enable, 18

F

file allocation
  defined, 116
  list policies, 119
  policies, defined, 116
  precedence rules, 117
  set file and directory policies, 117
  set segment preferences, 118
  tune policy settings, 119
file serving nodes
  add a CIFS share, 52
  delete, 34
  modify a CIFS share, 54
  SAN I/O, 9
  segment management, 8
  unmount a file system, 16
  view CIFS shares, 50
file systems
  32-bit mode, disable, 32
  64-bit mode, 11
  allocation policy, 8
  check and repair, 34
  components of, 9
  create
    from CLI, 14
    from GUI, 11
    options for, 11
    tiered file system, 109
  delete, 33
  delete tier policy rule, 112
  disk space information, 29
  expand tiered file system, 110
  Export Control, enable, 18
  extend, 29
  file allocation, 116
  list information
    tier policy, 111
    tiers on file system, 111
  list+found directory, 28
  mount, 15, 16
  mountpoints, create from CLI, 16
  mountpoints, delete, 15, 16

mountpoints, view, 15, 16
NFS
export, 38
unexport, 39
quotas, 19
remote replication, 87
segments
assign to tiers, 111
defined, 8
rebalance, 30
remove from tiers, 111
snapshots, 96
structure of, 8
troubleshooting, 35
unmount, 16
view summary information, 26

H
help
obtaining, 120
hostgroups, 16
HP
technical support, 120
HP websites, 120
HTTP
authentication methods, 42
configuration, 75
configuration profile, 76
configure local groups, 45
configure local users, 47
HTTP Wizard, 76
share, access, 81
share, configure, 79
virtual host, 78

L
Linux static user mapping with Active Directory, 59
Linux X9000 clients
disk space information, 29
Local users and groups, 45
logical volumes
view information, 26
lost+found directory, 28

M
management console
local users requirement, 44
Microsoft Management Console
manage CIFS shares, 54
mounting, file system, 15, 16
mountpoints
create from CLI, 16
delete, 15, 16
view, 15, 16

N
NFS clients
autoconnect, 40
set up, 40
NFS file systems
autoconnect NFS clients, 40
configure NFS server threads, 38
export, 38
set up NFS clients, 40
unexport, 39

P
physical volumes
delete, 33
view information, 25

Q
quotas, file system
delete, 24
enable, 19
operation of, 19
set for users and groups, 19

R
rebalancing segments, 30
stop tasks, 32
track job progress, 32
view job status, 32
related documentation, 120
remote replication
configure, 89
continuous, 87
defined, 87
export target file system, 90
identify host and NIC preferences, 89
intra-cluster, 88
many-to-many, 88
many-to-one, 88
pause replication, 93
query replication tasks, 94
register clusters, 89
remote cluster, 88
resume replication, 93
run-once, 87
start intra-cluster task, 93
start remote cluster task, 93
start run-once task, 93
stop replication, 93

S
SegmentNotAvailable alert, 35
SegmentRejected alert, 36
segments
defined, 8
delete, 33
rebalance, 30
stop tasks, 32
track job progress, 32
view job status, 32
SMB server consolidation, 62
SMB signing, 51
snapshots
automated, 97

create from CLI, 101
create on GUI, 98
delete snapshot scheme, 103
modify snapshot scheme, 103
schedule from CLI, 102
view snapshot scheme from CLI, 103
clear invalid snapshot, 103
create, 103
defined, 96
delete, 103
discover LUNs, 97
list storage allocation, 97
mount, 103
register the snapshot partition, 97
set up the snapshot partition, 96
view information about, 104
Subscriber's Choice, HP, 120

T

technical support
HP, 120
service locator website, 120

U

unmounting, file systems, 16

V

volume groups
delete, 33
view information, 25

W

websites
HP Subscriber's Choice for Business, 120

X

X9000 clients
delete, 34
hostgroups, 16
locally mount a file system, 17
locally unmount file system, 17
unmount a file system, 16